arXiv:0904.1366v2 [cs.DB] 14 Apr 2009

A Unified Approach to Ranking in Probabilistic Databases

Jian Li Barna Saha Amol Deshpande
{lijian, barna, amol}@cs.umd.edu
University of Maryland at College Park

April 14, 2009

Abstract

The dramatic growth in the number of application domains that naturally generate probabilistic, uncertain data has resulted in a need for efficiently supporting complex querying and decision-making over such data. In this paper, we present a unified approach to ranking and top-k query processing in probabilistic databases by viewing it as a multi-criteria optimization problem, and by deriving a set of *features* that capture the key properties of a probabilistic dataset that dictate the ranked result. We contend that a single, specific ranking function may not suffice for probabilistic databases, and we instead propose two *parameterized ranking functions*, called PRF^ω and PRF^e, that generalize or can approximate many of the previously proposed ranking functions. We present novel *generating functions*-based algorithms for efficiently ranking large datasets according to these ranking functions, even if the datasets exhibit complex correlations modeled using *probabilistic and/xor trees* or *Markov networks*. We further propose that the parameters of the ranking function be learned from user preferences, and we develop an approach to learn those parameters. Finally, we present a comprehensive experimental study that illustrates the effectiveness of our parameterized ranking functions, especially PRF^e, at approximating other ranking functions and the scalability of our proposed algorithms for exact or approximate ranking.

1 Introduction

Recent years have seen a dramatic rise in the number of applications domains that naturally generate uncertain data and that demand aid for complex, decision-support queries over them. In part, this has been due to the increasing prevalence of applications such as information retrieval [16], data integration and cleaning [2, 12], text analytics [25, 21], social network analysis [1] etc. At the same time, large-scale instrumentation of nearly every aspect of our world using sensor monitoring infrastructures has resulted in an abundance of uncertain, noisy data [11, 6].

By their very nature, many of these applications require support for ranking over large scale datasets. For instance, consider a *House Search* application, where a user is searching for a house using a real estate sales dataset: *House(id, price, size, zipcode, . . .)*. However, such a dataset, which may be constructed by crawling and combining data from multiple sources, is inherently uncertain and noisy. In fact, the houses that the user prefers the most, are also the most likely to be sold by now. We may denote such uncertainty by associating with each advertisement a *probability* that it is still valid. However, incorporating such uncertainties into the returned answers is a challenge, considering the complex interplay between the score of a house by itself, and the probability that the advertisement is still valid.

The importance of this natural problem has led to much work on ranking in probabilistic databases in recent years. That prior work (which we review in more detail later) has proposed many different functions for combining the scores and the probabilities. We begin with a systematic exploration of these issues by recognizing that ranking in probabilistic databases is inherently a multi-criteria optimization problem, and by deriving a set of *features*, the key properties of a probabilistic dataset that influence the ranked result. We empirically illustrate the diverse and conflicting behavior of several natural ranking functions, and argue that a single specific ranking function may not be appropriate to rank different uncertain databases that we may encounter in practice. Furthermore, different users may weigh the features differently, resulting in different rankings over the same dataset. We then define a general and powerful ranking function, called *PRF*, that allows us to explore the space of possible ranking functions. We discuss its

relationship to previously proposed ranking functions, and also identify two specific parameterized ranking functions, called PRF^ω and PRF^e, as being interesting. The PRF^ω ranking function is essentially a linear, weighted ranking function that resembles the scoring functions typically used in information retrieval, web search etc. [22, 28, 5, 10]. We observe that such a function may not be suitable for ranking large-scale probabilistic databases due to its high running time, and instead propose PRF^e, which uses a single parameter, and can approximate previously proposed ranking functions for probabilistic databases very well.

We then develop novel algorithms based on *generating functions* to efficiently rank the tuples in a probabilistic dataset using any PRF ranking function. Our algorithm can handle a probabilistic dataset with arbitrary correlations; however, it is particularly efficient when the probabilistic database contains only *mutual exclusivity* and/or *mutual co-existence* correlations (called *probabilistic and/xor trees* [30]). Our results apply to some of the previously proposed ranking functions as well (one of our results was also independently obtained by Yi et al. [36]). Our main contributions can be summarized as follows:

- We develop a framework for learning ranking functions over probabilistic databases by identifying a set of key *features* and by proposing several parameterized ranking functions.
- We present novel algorithms based on *generating functions* that enable highly efficient processing of top-k queries over very large datasets. Our key algorithm is an $O(n \log(n))$ algorithm ($O(n)$ if the dataset is pre-sorted by score) for evaluating a PRF^e function over datasets with low correlations (more specifically, constant height probabilistic and/xor trees).
- We present a polynomial time algorithm for computing the top-k answers for a correlated dataset, where the correlations are represented using a bounded-treewidth Markov network. The algorithm we present is actually for computing the probability that a given tuple is ranked at a given position across all the possible worlds, and is of independent interest.
- We develop a novel, DFT-based algorithm for approximating an arbitrary weighted ranking function using a linear combination of PRF^e functions.
- We present a comprehensive experimental study over several real and synthetic datasets, comparing the behavior of the ranking functions and the effectiveness of our proposed algorithms.

Outline: We begin with a brief discussion of the related work (Section 2). In Section 3, we review our probabilistic database model and the prior work on ranking in probabilistic databases, and propose two parameterized ranking functions. In Section 4, we present our generating functions-based algorithms for ranking. We then present an approach to approximate different ranking functions using our parameterized ranking functions, and to learn a ranking function from user preferences (Section 5). We then briefly sketch an extension of our ranking algorithms to correlated datasets where the correlations are modeled using Markov networks (Section 6). We conclude with an experimental study in Section 7.

2 Related Work

There has been much work on managing probabilistic, uncertain, incomplete, and/or fuzzy data in database systems (see e.g. [29, 16, 6, 8, 35, 3]). With a rapid increase in the number of application domains where uncertain data arises naturally, such as data integration, information extraction, sensor networks, pervasive computing etc., this area has seen renewed interest in recent years [17]. This work has spanned a range of issues from theoretical development of data models and data languages, to practical implementation issues such as indexing techniques, and several research efforts are underway to build systems to manage uncertain data (e.g. MYSTIQ [8], Trio [35], ORION [6], MayBMS [3]). The approaches can be differentiated between *tuple-level uncertainty*, where "existence" probabilities are attached to the tuples of the database, and *attribute-level uncertainty*, where (possibly continuous) probability distributions are attached to the attributes. The proposed approaches differ further based on whether they consider correlations or not. Most work in probabilistic databases has either assumed independence [16, 8] or has restricted the correlations that can be modeled [29, 2, 32]. More recently, several approaches have been presented that allow representation of arbitrary correlations [33, 3].

Time	Car Loc	Plate No	Speed	...	Prob	Tuple Id
11:40	L1	X-123	120	...	0.4	t_1
11:55	L2	Y-245	130	...	0.7	t_2
11:35	L3	Y-245	80	...	0.3	t_3
12:10	L4	Z-541	95	...	0.4	t_4
12:25	L5	Z-541	110	...	0.6	t_5
12:15	L6	L-110	105	...	0.1	t_6

World	Prob
$pw_1 = \{t_1, t_2, t_4, t_6\}$	.112
$pw_2 = \{t_1, t_2, t_5, t_6\}$	.168
$pw_3 = \{t_1, t_3, t_4, t_6\}$	.048
$pw_4 = \{t_1, t_3, t_5, t_6\}$	.072
$pw_5 = \{t_2, t_4, t_6\}$	.168
$pw_6 = \{t_2, t_5, t_6\}$	.252
$pw_7 = \{t_3, t_4, t_6\}$	.072
$pw_8 = \{t_3, t_5, t_6\}$	.108

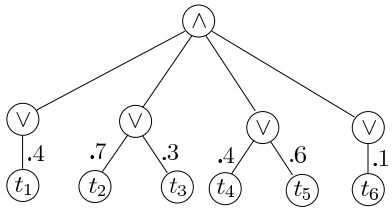

Figure 1: Example of a probabilistic database which contains automatically captured information about speeding cars. Tuple t_2 and t_3 (similarly, t_4 and t_5) are mutually exclusive. The corresponding and/xor tree compactly encodes these correlations.

The area of ranking and top-k query processing has also seen much work in databases (see Ilyas et al. [24] for a survey). More recently, several researchers have considered top-k query processing in probabilistic databases. Soliman et al. [34] defined the problem of ranking over probabilistic databases, and proposed two ranking functions to combine tuple scores and probabilities. Yi et al. [36] present improved algorithms for the same ranking functions. Zhang and Chomicki [37] present a desiderata for ranking functions. Ming Hua et al. [23] recently presented a different approach called *probabilistic threshold queries*. Finally, Cormode et al. [7] also present a semantics of ranking functions and a new ranking function called *expected rank*. We will review these ranking functions in detail in next section. There has also been work on top-k query processing in probabilistic databases where the ranking is by the result tuple *probabilities* (i.e., probability and score are identical) [31]. The main challenge in that work is efficient computation of the probabilities, whereas we assume that the probability and score are either given or can be computed easily.

3 Problem Formulation

We begin with defining our model of a probabilistic database, called *probabilistic and/xor tree* [30], that allows capturing several common types of correlations. We then review the prior work on top-k query processing in probabilistic databases, and argue that a single specific ranking function may not capture the intricacies of ranking with uncertainty. We then present our parameterized ranking functions, PRF^ω and PRF^e.

3.1 Probabilistic Database Model

We use the prevalent *possible worlds semantics* for probabilistic databases [8]. We denote a probabilistic relation with tuple uncertainty by D_T, where T denotes the set of tuples (see Section 4.4 for extensions of our algorithms to attribute uncertainty). The set of all possible worlds is denoted by $PW = \{pw_1, pw_2, ..., pw_n\}$. Each tuple $t_i \in T$ is associated with an existence probability $\mathsf{Pr}(t_i)$ and a score $\mathsf{score}(t_i)$, computed based on a scoring function $\mathsf{score} : T \rightarrow \mathbb{R}$. Usually $\mathsf{score}(t)$ is computed based on the tuple attribute values and measures the relative user preference for different tuples. In a deterministic database, tuples with higher scores should be ranked higher. We use $r_{pw} : T \rightarrow \{1, \ldots, n\} \cup \{\infty\}$ to denote the rank of the tuple t in a possible world pw according to score. If t does not appear in the possible world pw, we let $r_{pw}(t) = \infty$. We say t_1 *ranks higher* than t_2 in the possible world pw if $r_{pw}(t_1) < r_{pw}(t_2)$. For each tuple t, we define a random variable $r(t)$ which denotes the rank of t in D_T. In other words, $\mathsf{Pr}(r(t) = k)$ is the total probability of the possible worlds where t is ranked at position k.

Probabilistic And/Xor Tree Model: Our algorithms can handle arbitrarily correlated relations with correlations modeled using Markov networks (Section 6). However, in most of this paper, we focus on the *probabilistic and/xor tree model*, introduced in our prior work [30], that can capture only a more restricted set of correlations, but admits highly efficient ranking algorithms. More specifically, an and/xor tree captures two types of correlations: (1) *mutual exclusivity* (denoted ⩒ (*xor*)) and (2) *mutual co-existence* (⩓ (*and*)). Two events satisfy the mutual co-existence correlation if, in any possible world, either both events occur or neither occurs. Similarly two events are mutually exclusive if there is no possible world where both happens.

Definition 1 *A probabilistic and/xor tree $\mathcal{T}$ is a tree where each leaf is a singleton tuple and each inner node has a mark, ⩒ or ⩓. For each ⩒ node u and each of its children v, there is a nonnegative value $p(u, v)$ associated with*

the edge (u, v)*. Moreover, we require* $\sum_{v:(u,v)} p(u, v) \leq 1$*. Let* $\mathcal{T}_v$ *be the subtree rooted at* v *and* $v_1, \ldots, v_l$ *be* v*'s children. The subtree* $\mathcal{T}_v$ *inductively defines a random subset* S_v *of its leaves by the following independent process:*

- *If* v *is a leaf,* $S_v = \{v\}$.
- *If* v *is a* $\bigvee$ *node, then* $S_v = \begin{cases} S_{v_i} & \text{with prob. } p(v, v_i) \\ \emptyset & \text{otherwise} \end{cases}$
- v *is a* $\bigwedge$ *node, then* $S_v = \cup_i S_{v_i}$

x-tuples (which can be used to specify mutual exclusivity correlations between tuples) correspond to the special case where we have a tree of *height* 2, with a $\bigwedge$ node as the root and only $\bigvee$ nodes in the second level. Figure 1 shows an example of an and/xor tree that models the data from a traffic monitoring application [34], where the tuples represent automatically captured traffic data. For example, the leftmost $\bigvee$ node indicates t_1 is present with probability .4 and the second $\bigvee$ node dictates that exactly one of t_2 and t_3 should appear. The topmost $\bigwedge$ node tells us the random sets derived from these $\bigvee$ nodes coexist.

Probabilistic and/xor trees significantly generalize *x-tuples* [32, 36], block-independent disjoint tuples model, and p-or-sets [9], and can in fact represent a finite set of arbitrary possible worlds [30]. The correlations captured by such a tree can be represented by probabilistic c-tables [20] and provenance semirings [19]. However, that does not directly imply an efficient algorithm for ranking. *We note that no prior work on ranking in probabilistic databases has considered more complex correlations than x-tuples.*

3.2 Ranking over Probabilistic Data: Definitions and Prior Work

The interplay between probabilities and scores complicates the semantics of ranking in probabilistic databases. This was observed by Soliman et al. [34], who first considered this problem and presented two definitions of top-k queries in probabilistic databases. Several other definitions of ranking have been proposed since then. We briefly review the ranking functions we consider in this work.

- *Uncertain Top-*k *(*U-Top*)* [34]: Here the query returns the k-tuple set that appears as the top-k answer in most possible worlds (weighted by the probabilities of the worlds).
- *Uncertain Rank-*k *(*U-Rank*)* [34]: At each rank i, we return the tuple with the maximum probability of being at the i'th rank in all possible worlds. In other words, *U-Rank* returns: $\{t_i^*, i = 1, 2, .., \mathrm{k}\}$, where $t_i^* = argmax_t(\mathsf{Pr}(r(t) = i))$.
- *Probabilistic Threshold Top-k (*$PT(h)$*)* [23]: The original definition of a probabilistic threshold query asks for all tuples with probability of being in top-h answer larger than a pre-specified threshold, i.e., all tuples t such that $\mathsf{Pr}(r(t) \leq h) > threshold$. For consistency with other ranking definitions, we slightly modify the definition and instead ask for the k tuples with the largest $\mathsf{Pr}(r(t) \leq h)$ values.
- *Expected Ranks (*Exp-Rank*)* [7]: The tuples are ranked in the increasing order by: $\sum_{pw \in PW} \mathsf{Pr}(pw) r_{pw}(t)$, where $r_{pw}(t)$ is defined to be $|pw|$ if $t \notin pw$.
- *Expected Score (*E-Score*)*: Another natural ranking function, also considered by [7], is simply to rank the tuples by their expected score, $\mathsf{Pr}(t)\mathsf{score}(t)$.

Normalized Kendall Distance: To compare different ranking functions or criteria, we need a distance measure to evaluate the closeness of two top-k answers. We use the prevalent *Kendall tau* distance defined for comparing top-k answers for this purpose [14]. It is also called *Kemeny distance* in the literature and is considered to have many advantages over other distance metrics [13]. Let $\mathcal{R}_1$ and $\mathcal{R}_2$ denote two full ranked lists, and let $\mathcal{K}_1$ and $\mathcal{K}_2$ denote the top-k ranked tuples in $\mathcal{R}_1$ and $\mathcal{R}_2$ respectively. Then *Kendall tau distance* between $\mathcal{K}_1$ and $\mathcal{K}_2$ is defined to be: $\mathsf{dis}(\mathcal{K}_1, \mathcal{K}_2) = \sum_{(i,j) \in P(\mathcal{K}_1, \mathcal{K}_2)} \hat{K}(i, j)$, where $P(\mathcal{K}_1, \mathcal{K}_2)$ is the set of all unordered pairs of $\mathcal{K}_1 \cup \mathcal{K}_2$; $\hat{K}(i, j) = 1$ if it can be inferred from $\mathcal{K}_1$ and $\mathcal{K}_2$ that i and j appear in opposite order in the two full ranked lists $\mathcal{R}_1$ and $\mathcal{R}_2$, otherwise $\hat{K}(i, j) = 0$. Intuitively the Kendall distance measures the number of inversions or flips between the two rankings. For ease of comparison, we divide the Kendall distance by k^2 to obtain *normalized Kendall* distance, which always lies in $[0, 1]$.

A higher value of the Kendall distance indicates a larger disagreement between the two top-k lists. It is easy to see that if the Kendall distance between two top-k answers is δ, then the two answers must share at least $1 - \sqrt{\delta}$ fraction of tuples (so if the distance is 0.09, then the top-k answers share at least 70%, and typically 90% or more tuples). The distance is zero only if two top-k answers are completely disjoint.

	E-Score	$PT(100)$	*U-Rank*	*Exp-Rank*	*U-Top*
E-Score	–	0.1241	0.3027	0.7992	0.2760
$PT(100)$	0.1241	–	0.3324	0.9290	0.3674
U-Rank	0.3027	0.3324	–	0.9293	0.2046
Exp-Rank	0.7992	0.9290	0.9293	–	0.9456
U-Top	0.2760	0.3674	0.2046	0.9456	–

IIP-100,000 ($\mathrm{k} = 100$**)**

	E-Score	$PT(100)$	*U-Rank*	*Exp-Rank*	*U-Top*
E-Score	–	0.8642	0.8902	0.0044	0.9258
$PT(100)$	0.8642	–	0.3950	0.8647	0.5791
U-Rank	0.8902	0.3950	–	0.8907	0.3160
Exp-Rank	0.0044	0.8647	0.8907	–	0.9263
U-Top	0.9258	0.5791	0.3160	0.9263	–

Syn-IND Dataset with 100,00 tuples ($\mathrm{k} = 100$**)**

Table 1: Normalized Kendall distance among various ranking functions for two datasets

Comparing Ranking Functions: We compared the top-100 answers returned by the five ranking functions with each other using the normalized Kendall distance, for two datasets with 100,000 independent tuples each (see Section 7 for a description of the datasets). Table 1 shows the results of this experiment. As we can see, the five ranking functions return wildly different top-k answers for the two datasets, with no obvious trends. For the first dataset, *Exp-Rank* behaves very differently from all other functions, whereas for the second dataset, *Exp-Rank* happens to be quite close to *E-Score*. However both of them deviate largely from *U-Top*, $PT(h)$ and *U-Rank*. The behavior of *E-Score* is very sensitive to the dataset, especially the score distribution: it is close to $PT(h)$ and *U-Rank* for the first dataset, but far away from all of them in the second dataset (by looking into the results, it shares less than 15 tuples with the Top-100 answers of the others). We observed similar behavior for other datasets, and for datasets with correlations.

This simple experiment illustrates the issues with ranking in probabilistic databases – although several of these definitions seem natural, the wildly different answers they return indicate that none of them could be the "right" definition.

We also observe that in large datasets, *Exp-Rank* tends to give very high priority to a tuple with a high probability even if it has a low score. In our synthetic dataset Syn-IND-100,000 with expected size ≈ 50000, t_2 (the tuple with 2nd highest score) has probability approximately 0.98 and t_{1000} (the tuple with 1000th highest score) has probability 0.99. The expected ranks of t_2 and t_{1000} are approximately 10000 and 6000 respectively, and hence t_{1000} is ranked above t_2 even though t_{1000} is only slightly more probable.

3.3 Parameterized Ranking Functions

Ranking in uncertain databases is inherently a multi-criteria optimization problem, and it is not always clear how to rank two tuples that dominate each other along different axes. Consider a database with two tuples t_1 (score = 100, $\mathsf{Pr}(t_1) = 0.5$), and t_2 (score = 50, $\mathsf{Pr}(t_2) = 1.0$). Even in this simple case, it is not clear whether to rank t_1 above t_2 or vice versa. This is an instance of the classic risk-reward trade-off, and the choice between these two options largely depends on the application domain and/or user preferences.

We propose to follow the traditional approach to dealing with such tradeoffs, by identifying a set of *features*, by defining a parameterized ranking function over these features, and by learning the parameters (weights) themselves using user preferences [22, 28, 5, 10]. To achieve this, we propose a family of ranking functions, parameterized by one or more parameters, and design algorithms to efficiently find the top-k answer according to any ranking function from these families. Our general ranking function, PRF, directly subsumes some of the previously proposed ranking

functions, and can also be used to approximate other ranking functions. Moreover, the parameters can be learned from user preferences, which allows us to adapt to different scenarios and different application domains.

Features: Although it is tempting to use the tuple probability and the tuple score as the features, a ranking function based on just those two will be highly sensitive to the actual values of the scores; further, such a ranking function will be insensitive to the correlations in the database, and hence cannot capture the rich interactions between ranking and possible worlds.

Instead we propose to use the following set of features: for each tuple t, we have n features, $\mathsf{Pr}(r(t) = i), i = 1, \cdots, n$, where n is the number of tuples in the database. This set of features succinctly captures the possible worlds. Further, correlations among tuples, if any, are naturally accounted for when computing the features. We note that in most cases, we do not explicitly compute all the features, and instead design algorithms that can directly compute the value of the overall ranking function.

Ranking Functions: Next we define a general ranking function which allows exploring the trade-offs discussed above.

Definition 2 *Let $\omega : T \times \mathbb{N} \rightarrow \mathbb{C}$ be a* weight function, *where $\mathbb{C}$ is the set of complex numbers. The* parameterized ranking function (PRF), *$\Upsilon_\omega : T \rightarrow \mathbb{C}$ in its most general form is defined to be:*

$$\begin{aligned}\Upsilon_\omega(t) &= \sum_{pw:t\in pw} \omega(t, r_{pw}(t)) \cdot \mathsf{Pr}(pw) \\ &= \sum_{pw:t\in pw} \sum_{i>0} \omega(t,i)\mathsf{Pr}(pw \wedge r_{pw}(t) = i) \\ &= \sum_{i>0} \omega(t,i) \cdot \mathsf{Pr}(r(t) = i).\end{aligned}$$

A top-k query returns k *tuples with the highest $|\Upsilon_\omega|$ values.*

In most cases, ω is a real positive function and we just need to find the k tuples with highest Υ_ω values. However we allow ω to be a complex function in order to approximate other functions efficiently (see Section 5.1). Depending on the actual function ω, we get different ranking functions with diverse behaviors. We illustrate some of these choices and relate them to prior ranking functions[1]. We omit the subscript ω if the context is clear.

- If $\omega(t, i) = 1$, the result is the set of k tuples with the highest probabilities [31].
- By setting $\omega(t, i) = \mathsf{score}(t)$, we get the *E-Score* :
  $$\Upsilon(t) = \textstyle\sum_{pw:t\in pw} \mathsf{score}(t)\mathsf{Pr}(pw) = \mathsf{score}(t)\mathsf{Pr}(t) = \mathsf{E}(\mathsf{score}(t))$$
- $PRF^\omega(h)$: One important class of ranking functions is when $\omega(t, i) = w_i$ (i.e., independent of t) and $w_i = 0 \forall i > h$ for some positive integer h. This forms one of prevalent classes of ranking functions used in domains such as information retrieval and machine learning, with the weights typically learned from user preferences [22, 28, 5, 10].
- Two special cases of the PRF^ω function are:
  1. $\omega(i) = \begin{cases} 1, & i \leq h \\ 0, & \text{otherwise} \end{cases}$. If we return k tuples with highest $\Upsilon_\omega(t)$ value, we have exactly the answer for $PT(h)$.
  2. $\omega_j(i) = \begin{cases} 1, & \text{i=j} \\ 0, & \text{otherwise} \end{cases}$ for some $1 \leq j \leq \mathrm{k}$. We can see the tuple with largest Υ_{ω_j} value is the rank-j answer in *U-Rank* query [34].
     This allows us to compute the *U-Rank* answer by evaluating $\Upsilon_{\omega_j}(t)$ for all $t \in T$ and $j = 1, \ldots, \mathrm{k}$.

[1]The definition of the *U-Top* introduced in [34] requires the retrieved k tuples belongs to a valid possible world. However, it is not required in our definition, and hence it is not possible to simulate *U-Top* using PRF.

- $\underline{PRF^e(\alpha)}$: Finally, we define PRF^e to be a special case of the PRF^ω function, where $w_i = \omega(i) = \alpha^i$, where α is a constant and may be a real or a complex number.

PRF^ω and PRF^e form the two parameterized ranking functions that we propose in this work. Although PRF^ω is the more natural ranking function and has been used elsewhere, PRF^e is more suitable for ranking in probabilistic databases for various reasons. First, the features as we have defined above are not completely arbitrary, and the features $\mathsf{Pr}(r(t) = i)$ for small i are clearly more important than the ones for large i. Hence in most cases we would like the weight function, $\omega(i)$, to be monotonically non-increasing. PRF^e naturally captures this behavior (as long as $|\alpha| < 1$). More importantly, we can compute the PRF^e function in $O(n \log(n))$ time ($O(n)$ time if the dataset is pre-sorted by score) even for datasets with low degrees of correlations (i.e., modeled by and/xor trees with low heights). This makes it significantly more attractive for ranking over large datasets.

Furthermore, ranking by $PRF^e(\alpha)$, with suitably chosen α, can approximate rankings by many other functions reasonably well even with only real α. Finally, a linear combination of exponential functions, with complex bases, is known to be very expressive in representing other functions [4]. We make use of this fact to approximate many ranking functions by linear combinations of a small number of PRF^e functions, thus significantly speeding up the running time. We revisit this in Section 5.1.

4 Ranking Algorithms

We next present an algorithm for efficiently ranking according to a PRF function. We first present the basic idea behind our algorithm assuming mutual independence, and then consider correlated tuples with correlations represented using an and/xor tree. We then present a very efficient algorithm for ranking using a PRF^e function, and then briefly discuss how to handle attribute uncertainty.

4.1 Assuming Tuple Independence

First we show how the PRF function can be computed in $O(n^2)$ time for general weight function ω, and for a given set of tuples $T = \{t_1, \ldots, t_n\}$. In all our algorithms, we assume that $\omega(t, i)$ can be computed in $O(1)$ time. Clearly it would be sufficient to compute $\mathsf{Pr}(r(t) = j)$ for any tuple t and $1 \leq j \leq n$ in $O(n^2)$ time. Given these values, we can directly compute the values of $\Upsilon(t)$ in $O(n^2)$ time.

We first sort the tuples in a non-increasing order by their score function; assume $t_1, \ldots, t_n$ indicates this sorted order. Suppose now we want to compute $\mathsf{Pr}(r(t_i) = j)$. Let $T_i = \{t_1, t_2, \ldots, t_i\}$ and σ_i be an indicator variable that takes value 1 if t_i is present, and 0 if otherwise. Further, let $\sigma = \langle \sigma_1, \ldots, \sigma_n \rangle$ denote a vector containing all the indicator variables. Then, we can write $\mathsf{Pr}(r(t_i) = j)$ as follows:

$$\begin{aligned} \mathsf{Pr}(r(t_i) = j) &= \mathsf{Pr}(t_i) \sum_{pw : |pw \cap T_{i-1}| = j-1} \mathsf{Pr}(pw) \\ &= \mathsf{Pr}(t_i) \sum_{\sigma : \sum_{l=1}^{i-1} \sigma_l = j-1} \prod_{l<i:\sigma_l=1} \mathsf{Pr}(t_l) \prod_{l<i:\sigma_l=0} (1 - \mathsf{Pr}(t_l)) \end{aligned}$$

The first equality says that tuple t_i ranks at the jth position if and only t_i and exactly $j-1$ tuples from T_{i-1} are present in the possible world. The second equality holds since the tuples are independent of each other. The naive method to evaluate the above formula by explicitly listing all possible worlds needs exponential time. Now, we turn to the generating function method based on which a polynomial time algorithm can be obtained. Recall the coefficient of x^k in $\mathcal{F}(x) = \prod_{i=1}^{n}(a_i + b_i x)$ is $\sum_{|\beta|=k} \prod_{i:\beta_i=0} a_i \prod_{i:\beta_i=1} b_i$ where $\beta = \langle \beta_1, \ldots, \beta_n \rangle$ and $|\beta|$ denote the number of 1's in β [18]. Now consider the following generating function:

$$\begin{aligned} \mathcal{F}^i(x) &= \left(\prod_{t \in T_{i-1}} \Big(1 - \mathsf{Pr}(t) + \mathsf{Pr}(t) \cdot x\Big) \right) (\mathsf{Pr}(t_i) \cdot x) \\ &= \sum_{j \geq 0} c_j x^j. \end{aligned}$$

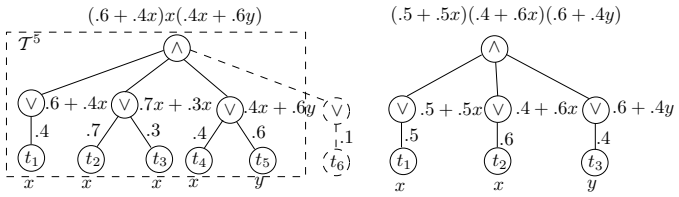

Figure 2: PRF computation on and/xor trees: the left figure corresponds to the database in Figure 1, whereas the right figure is the and/xor tree representation of the independent tuples in Example 1.

We can see that the coefficient c_j of x^j in the expansion of $\mathcal{F}^i$ is exactly the probability that t_i is at rank j, i.e., $\mathsf{Pr}(r(t_i) = j)$. We note $\mathcal{F}^i$ contains at most $i+1$ nonzero terms. Expanding $\mathcal{F}^i$ can be easily done in $O(i^2)$ time. This allows us to compute $\mathsf{Pr}(r(t_i) = j)$ for t_i in $O(i^2)$ time; $\Upsilon(t_i)$, in turn, can be written as:

$$\Upsilon(t_i) = \sum_j \omega(j) \cdot \mathsf{Pr}(r(t_i) = j) = \sum_j \omega(j) c_j \tag{1}$$

which can be computed in $O(i^2)$ time.

Example 1 *Consider a relation with 3 independent tuples t_1, t_2, t_3 (already sorted according to the score function) with existence probabilities $0.5, 0.6, 0.4$, respectively. The generating function for t_3 is:*

$$\mathcal{F}^3(x) = (.5 + .5x)(.4 + .6x)(.4x) = .2x + .5x^2 + .3x^3$$

This gives us:

$$\mathsf{Pr}(r(t_3) = 1) = .2, \mathsf{Pr}(r(t_3) = 2) = .5, \mathsf{Pr}(r(t_3) = 3) = .3$$

If we expand each $\mathcal{F}^i$ for $1 \le i \le n$ from scratch, we need $O(n^2)$ time for each $\mathcal{F}^i$ and $O(n^3)$ time in total. However, the expansion of $\mathcal{F}^i$ can be obtained from the expansion of $\mathcal{F}^{i-1}$ in $O(i)$ time as follows:

$$\mathcal{F}^i(x) = \frac{\mathsf{Pr}(t_i)}{\mathsf{Pr}(t_{i-1})} \mathcal{F}^{i-1}(x) \Big(1 - \mathsf{Pr}(t_{i-1}) + \mathsf{Pr}(t_{i-1})x\Big). \tag{2}$$

This trick gives us a $O(n^2)$ time complexity. See Algorithm 1 for the pseudocode. Note that $O(n^2)$ time is asymptotically optimal in general since the computation involves at least $O(n^2)$ probabilities, namely $\mathsf{Pr}(r(t_i) = j)$ for all $1 \le i, j \le n$.

Algorithm 1: IND-PRF-RANK(D_T)

```
F^0(x) = 1;
for i = 1 to n do
    F^i(x) = (Pr(t_i)/Pr(t_{i-1})) F^{i-1}(x) (1 - Pr(t_{i-1}) + Pr(t_{i-1})x) ;
    Expand F^i(x) in the form of sum_j c_j x^j ;
    Υ(t_i) = sum_{j=1}^{n} ω(t_i, j) c_j ;
return k tuples with largest Υ values;
```

For some specific ω functions, we may be able to achieve faster running time. For $PRF^\omega(h)$ functions, we only need to expand all $\mathcal{F}^i$'s up to x^h term since $\omega(i) = 0$ for $i > h$. Then, the expansion from $\mathcal{F}^{i-1}(x)$ to $\mathcal{F}^i(x)$ only takes $O(h)$ time. This yields an $O(n \cdot h + n \log(n))$ time algorithm. We note the above technique also gives an $O(n\mathrm{k} + n \log(n))$ time algorithm for answering the *U-Rank* top-k query (all the needed probabilities can be computed in that time), thus matching the best known upper bound by Yi et al. [36] (the original algorithm in [34] runs in $O(n^2\mathrm{k})$ time).

We remark that the generating function technique can be seen as a variant of dynamic programming in some sense; however, using it explicitly in place of the obscure recursion formula gives us a much cleaner view and allow us to generalize it to handle more complicated tuple correlations. This also leads to an algorithm for extremely efficient evaluation of PRF^e functions (Section 4.3).

4.2 Probabilistic And/Xor Trees

Next we generalize our algorithm to handle a correlated database where the correlations can be captured using an *and/xor* tree. As before, let $T = \{t_1, t_2, \ldots, t_n\}$ denote the tuples sorted in an non-increasing order of their score function, and let $T_i = \{t_1, t_2, \ldots, t_i\}$. Let $\mathcal{T}$ denote the and/xor tree that models the correlations.

Suppose now we need to compute $\mathsf{Pr}(r(t_i) = j)$. Since we don't care about any tuple with a smaller score, it suffices to consider only $\mathcal{T}_i$, the subtree of $\mathcal{T}$ induced by the leaf set T_i (namely, the union of all root-leaf paths with all leaves in T_i). Let $Ch(v) = \{v_1, \ldots, v_l\}$ denote the set of v's children. Let $p_v = \sum_{v_h \in Ch(v)} p(v, v_h)$. For each node $v \in \mathcal{T}_i$, we define generating function $\mathcal{F}_v^i(x, y)$ inductively as follows:

- If v is a leaf, $\mathcal{F}_v^i(x,y) = \begin{cases} x, & v \in T_i \setminus \{t_i\}; \\ y, & v = t_i. \end{cases}$.
- If v is a $\circledvee$ node,
  $\mathcal{F}_v^i(x,y) = (1 - p_v) + \sum_{v_h \in Ch(v)} \mathcal{F}_{v_h}^i(x,y) \cdot p(v, v_h)$
- If v is a $\circledwedge$ node, $\mathcal{F}_v^i(x,y) = \prod_{v_h \in Ch(v)} \mathcal{F}_{v_h}^i(x,y)$.

The generating function $\mathcal{F}^i$ for $\mathcal{T}_i$ is the generating function of its root. The following theorem [30] states the close relationship between the probabilities $\mathsf{Pr}(r(t_i) = j)$ and the coefficients of $\mathcal{F}^i$.

Theorem 1 *[30] Let c_j be the coefficient of the term $x^{j-1}y$ in the generating function $\mathcal{F}^i(x,y)$ defined as above. We have that: $\mathsf{Pr}(r(t_i) = j) = c_j$.*

Proof: Suppose $\mathcal{T}^i$ is rooted at r, $r_1, \ldots, r_h$ are r's children and $\mathcal{T}_l^i$ is the subtree rooted at r_l. W.l.o.g, we assume $t_i \in \mathcal{T}_h^i$ We denote by S^i (S_l^i) the random set of leaves generated according to model $\mathcal{T}^i$ ($\mathcal{T}_l^i$). We write $\mathcal{F}^i(x,y) = \sum_j c'_j x^j + (\sum_j c_j x^j)y$ We prove by induction on the height of the and/xor tree the following claims: $\mathsf{Pr}(|S^i| = j \wedge t_i \in S^i)) = c_j$ and $\mathsf{Pr}(|S^i| = j \wedge t_i \notin S^i)) = c'_j$ We consider two case. If r is a $\circledwedge$ node, we know from Definition 1 $S^i = \cup_{l=1}^h S_l^i$. Therefore,

$$\mathsf{Pr}(|S^i| = j \wedge t_i \in S^i) = \mathsf{Pr}(\sum_{l=1}^h |S_l^i| = j \wedge t_i \in S^i)$$

$$= \sum_{\sum_{l=1}^h j_l = j} \left(\prod_{l=1}^h \mathsf{Pr}(|S_l^i| = j_l) \right) \mathsf{Pr}(|S_h^i| = j_h \wedge t_i \in S_h^i)$$

Let $\mathcal{F}_l^i = \sum_j c'_{lj} x^j + (\sum_j c_{lj} x^j)y$ be the polynomial defined by $\mathcal{T}_l^i$. Now let us consider the coefficient c_j of the term $x^{j-1}y$ in $\mathcal{F}^i(x,y) = \prod_{l=1}^h \mathcal{F}_l^i(x,y)$. From the Algorithm 2, we know the degree of y in $\mathcal{F}_l^i$ is 0 for $1 \le l \le h-1$. So only the terms in $\mathcal{F}_h^i$ can contribute to the y terms in $\mathcal{F}^i$. Therefore, $c_j = \sum_{\sum_{l=1}^h j_l = j} \left(\prod_{l=1}^{h-1} c'_{lj_l} \right) c_{lj_h}$. Then it is easy to see, from induction hypothesis, the theorem follows. The other case when r is a $\circledvee$ node can be proved by simply observing $\mathsf{Pr}(|S^i| = j \wedge t_i \in S^i) = \mathsf{Pr}(|S_l^i| = j \wedge t_i \in S_l^i) p(r, r_h)$ and the coefficient c_j of the term $x^j y$ in $\mathcal{F}^i$ is $p(r, r_h) c_{hj}$. □

Example 2 *The generating function $\mathcal{F}^5$ for the left hand side tree in Figure 2 is $(.6 + .4x)x(.4x + .6y) = .24x^2 + .16x^3 + .36xy + .24x^2y$. So we get that $\mathsf{Pr}(r(t_5) = 3) = .24$. From Figure 1, we can also see $\mathsf{Pr}(r(t_5) = 3) = \mathsf{Pr}(pw_2) + \mathsf{Pr}(pw_4) = .24$. The right hand side of Figure 2 shows the probabilistic and/xor tree and the generating function computation for Example 1.*

Algorithm 2: ANDOR-PRF-RANK($\mathcal{T}$)

$\mathcal{T}^0 = \emptyset$;
for *i=1 to n* **do**
  $\mathcal{T}_i = \mathcal{T}_{i-1} \cup$ the path from t_i to the root;
  $\mathcal{F}^i(x, y) = GENE(\mathcal{T}_i, t_i)$;
  Expand $\mathcal{F}^i(x, y)$ in the form $\sum_j c'_j x^j + (\sum_j c_j x^{j-1})y$;
  $\Upsilon(t_i) = \sum_{j=1}^{n} \omega(t_i, j) c_j$;
return k tuples with largest Υ values;
Subroutine: $GENE(\mathcal{T}, t)$;
if $\mathcal{T}$ *is a singleton node* **then**
  if $\mathcal{T} = \{t\}$ **then return** y **else return** x
else
  $\mathcal{T}_i$ is the subtree rooted at r_i for $r_i \in Ch(r)$;
  $p = \sum_{r_i \in Ch(r)} p(r, r_i)$;
  if r *is a* $\otimes$ *node* **then**
    return $1 - p + \sum_{r_i \in Ch(r)} p(r, r_i) \cdot GENE(\mathcal{T}_i, t)$;
  if r *is a* $\oplus$ *node* **then**
    return $\prod_{r_i \in Ch(r)} GENE(\mathcal{T}_i, t)$;

See Algorithm 2 for the pseudocode of the algorithm.

If we expand $\mathcal{F}_v^i$ for each internal node v in a naive way (i.e., we do polynomial multiplication one by one), we can show the running time is $O(n^2)$ at each internal node and thus $O(n^3)$ overall. If we do divide-and-conquer at each internal node and apply FFT (Fast Fourier Transformation) for the multiplication of polynomials, the running time can be improved to $O(n^2 \log^2 n)$. The details can be found in Appendix A.

4.3 Computing a PRF^e Function

Next we present an $O(n \log(n))$ algorithm to evaluate a PRF^e function (the algorithm runs in linear time if the dataset is pre-sorted by score). If $\omega(i) = \alpha^i$, then we observe that:

$$\Upsilon(t_i) = \sum_{j=1}^{n} \mathsf{Pr}(r(t_i) = j)\alpha^j = \mathcal{F}^i(\alpha) \tag{3}$$

This surprisingly simple relationship suggests we don't have to expand the polynomials $\mathcal{F}^i(x)$ at all; instead we can evaluate the numerical value of $\mathcal{F}^i(\alpha)$ directly. Again, we note that the value $\mathcal{F}^i(\alpha)$ can be computed from the value of $\mathcal{F}^{i-1}(\alpha)$ in $O(1)$ time using Equation (2). Thus, we have $O(n)$ time algorithm to compute $\Upsilon(t_i)$ for all $1 \le i \le n$ if the tuples are pre-sorted.

Example 3 *Consider Example 1 and the PRF^e function for t_3. We choose $\omega(i) = .6^i$. Then, we can see that* $\mathcal{F}^3(x) = (.5 + .5x)(.4 + .6x)(.4x)$. *So,* $\Upsilon(t_3) = \mathcal{F}^3(.6) = (.5 + .5 \times .6)(.4 + .6 \times .6)(.4 \times .6) = .14592$.

We can use a similar idea to speed up the computation if the tuples are correlated and the correlations are represented using an and/xor tree. Suppose the generating function for $\mathcal{T}_i$ is $\mathcal{F}^i(x, y) = \sum_j c'_j x^j + (\sum_j c_j x^{j-1})y$ and $\Upsilon(t_i) = \sum_{j=1}^{n} \alpha^j c_j$. We observe an intriguing relationship between the PRF^e value and the generating function:

$$\begin{aligned} \Upsilon(t_i) &= \sum_j c_j \alpha^j = \Big(\sum_j c'_j \alpha^j + (\sum_j c_j \alpha^{j-1})\alpha\Big) - \sum_j c'_j \alpha^j \\ &= \mathcal{F}^i(\alpha, \alpha) - \mathcal{F}^i(\alpha, 0). \end{aligned}$$

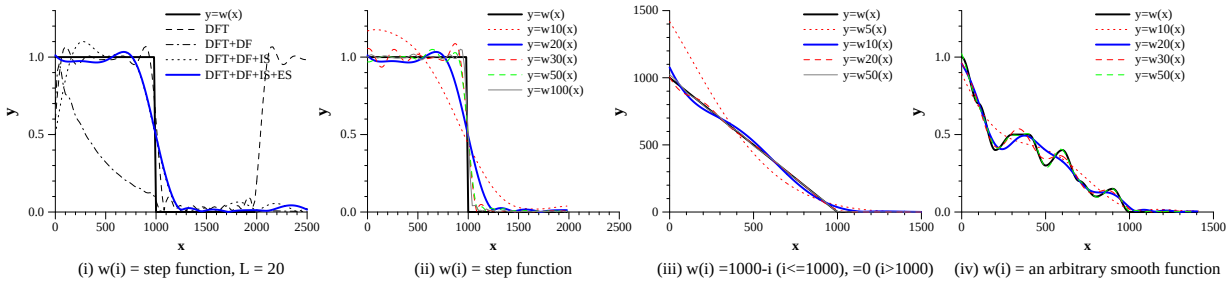

(i) w(i) = step function, L = 20 (ii) w(i) = step function (iii) w(i) =1000-i (i<=1000), =0 (i>1000) (iv) w(i) = an arbitrary smooth function

Figure 3: Approximating functions using linear combinations of complex exponentials

Given this, $\Upsilon(t_i)$ can be computed in linear time by bottom up evaluation of $\mathcal{F}^i(\alpha,\alpha)$ and $\mathcal{F}^i(\alpha,0)$ in $\mathcal{T}^i$. If we simply repeat it n times, once for each t_i, this gives us a $O(n^2)$ total running time.

By carefully sharing the intermediate results among computations of $\Upsilon(t_i)$, we can improve the running time to $O(n\log(n)+nd)$ where d is the height of the and/xor tree. We sketch this algorithm, which runs in iterations. Suppose the tuples are already pre-sorted by their scores. In iteration i, leaf t_i (the i'th tuple in score order) is added to the tree and the computations are done along the path from t_i to the root. Specifically, the algorithm maintains the following information in each inner node v: the numerical value of $\mathcal{F}^i_v(\alpha,\alpha)$ and $\mathcal{F}^i_v(\alpha,0)$. The values on node v need to be updated when the value of one of its children changes. Therefore, in each iteration, the computation only happens on the path from t_i to the root. Since we update at most d nodes for each newly added node, the running time is $O(nd)$. The updating rule for $\mathcal{F}^i_v(.,.)$(both $\mathcal{F}^i_v(\alpha,\alpha)$ and $\mathcal{F}^i_v(\alpha,0)$) in node v is as follows. We assume v's child, say u, just had its values changed.

1. v is a $\wedge$ node, $\mathcal{F}^i_v(.,.) \leftarrow \mathcal{F}^{i-1}_v(.,.)\mathcal{F}^i_u(.,.)/\mathcal{F}^{i-1}_u(.,.)$
2. v is a $\vee$ node, then:
   $\mathcal{F}^i_v(.,.) \leftarrow \mathcal{F}^{i-1}_v(.,.) + p(v,u)\mathcal{F}^i_u(.,.) - p(v,u)\mathcal{F}^{i-1}_u(.,.)$

We note that, for the case of *x-tuples*, which can be represented using a two-level tree, this gives us an $O(n\log(n))$ algorithm for ranking according to PRF^e.

4.4 Attribute Uncertainty or Uncertain Scores

We briefly sketch how we can do ranking over tuples with discrete attribute uncertainty where the uncertain attributes are part of the tuple scoring function (if the uncertain attributes do not affect the tuple score, then they can be ignored for the ranking purposes). More generally, this approach can handle the case when there is a discrete probability distribution over the score of the tuple.

The algorithm works by treating the alternatives of the tuples (with a separate alternative for each different possible score for the tuple) as different tuples, and adding an *xor* constraint over the alternatives. We can then use the algorithm for the probabilistic and/xor tree model to find the values of the PRF function for each resulting tuple separately. In a final step, we calculate the Υ score for each original tuple by adding the Υ scores of its alternatives If the original tuples were independent, the complexity of this algorithm is $O(n^2)$ for computing the PRF function, and $O(n\log(n))$ for computing the PRF^e function where n is the size of the input, i.e, the total number of different possible scores.

5 Approximating and Learning Ranking Functions

In this section, we discuss how to choose the PRF functions and their parameters. Depending on the application domain and the scenarios, there are two approaches to this: (1) If we know the ranking function we would like to use (say $PT(h)$), then we can either simulate or approximate it using appropriate PRF functions; (2) If we are instead provided user preferences data, we can learn the parameters from them. Clearly, we would prefer to use a PRF^e function, if possible, since it admits highly efficient ranking algorithms. For this purpose, we begin with presenting

an algorithm to find an approximation to an arbitrary PRF^ω function using a linear combination of PRF^e functions. We then discuss how to learn a PRF^ω function from user preferences, and finally present an algorithm for learning a single PRF^e function.

5.1 Approximating PRF^ω using PRF^e Functions

A linear combination of complex exponential functions is known to be very expressive, and can approximate many other functions very well [4]. Specifically, given a PRF^ω function, if we can write $\omega(i)$ as: $\omega(i) \approx \sum_{l=1}^{L} u_l \alpha_l^i$, then we have that:

$$\Upsilon(t) = \sum_i \omega(i)\mathsf{Pr}(r(t) = i) \approx \sum_{l=1}^{L} u_l \left(\sum_i \alpha_l^i \mathsf{Pr}(r(t) = i) \right)$$

This reduces the computation of $\Upsilon(t)$ to L individual PRF^e function computations, each of which only takes linear time. This gives us an $O(n\log(n) + nL)$ time algorithm for approximately ranking using PRF^ω function for independent tuples (as opposed to $O(n^2)$ for exact ranking).

Several techniques have been proposed for finding such approximations using complex exponentials [26, 4]. Those techniques are however computationally inefficient, involving computation of the inverses of large matrices and the roots of polynomials of high orders, and may be numerically unstable.

In this section, we present a clean and efficient algorithm, based on Discrete Fourier Transforms (DFT), for approximating a function $\omega(i)$, that approaches to zero for large values of i As we noted earlier, this captures the typical behavior of the $\omega(i)$ function. An example of such a function is the step function ($\omega(i) = 1 \forall i \leq h, = 0 \forall i > h$) which corresponds to the ranking function $PT(h)$. At a high level, our algorithm starts with a DFT approximation of $\omega(i)$ and then adapts it by adding several damping, scaling and shifting factors.

Discrete Fourier transformation (DFT) is a well known technique for representing a function as a linear combination of complex exponentials (also called *frequency domain representation*). More specifically, a discrete function $\omega(i)$ defined on a finite domain $[0, N-1]$ can be decomposed into exactly N exponentials as:

$$\omega(i) = \frac{1}{N} \sum_{k=0}^{N-1} \psi(k) e^{\frac{2\pi\jmath}{N} ki} \qquad i = 0, \ldots, N-1. \tag{4}$$

where $\psi(0), \cdots, \psi(N-1)$ denotes the DFT transform of $\omega(0), \cdots, \omega(N-1)$. If we want to approximate ω by fewer, say L, exponentials, we can instead use the L DFT coefficients with maximum absolute value. For clarity, we assume that $\psi(0), \ldots, \psi(L-1)$ are those coefficients. Then our approximation $\tilde{\omega}_L^{DFT}$ of ω by L exponentials is given by:

$$\tilde{\omega}_L^{DFT}(i) = \frac{1}{N} \sum_{k=0}^{L-1} \psi(k) e^{\frac{2\pi\jmath}{N} ki} \qquad i = 0, \ldots, N-1. \tag{5}$$

However, DFT utilizes only complex exponentials of unit norm, i.e., $e^{\jmath r}$ (where r is a real), which makes this approximation periodic (with a period of N). If we make N sufficiently large, say larger than the total number of tuples, then we usually need a large number of exponentials (L) to get a reasonable approximation. Moreover, computing DFT for very large N is computationally non-trivial. Furthermore, the number of tuples n may not be known in advance.

We next present a set of nontrivial tricks to adapt the base DFT approximation to overcome these shortcomings. To illustrate our method, we use the step function $\omega(i) = \begin{cases} 1, & i < N \\ 0, & i \geq N \end{cases}$ as our running example to show our method and the specific shortcoming it addresses. We assume $\omega(i)$ takes none-zero values within interval $[0, N-1]$ and the absolute values of both $\omega(i)$ and $\omega_L^{DFT}(i)$ are bounded by B.

1. (DFT) We perform pure DFT on the domain $[1, aN]$, where a is a small integer constant (typically < 10).

2. (Damping Factor (DF)) We introduce a damping factor $\eta \leq 1$ such that $B\eta^{aN} \leq \epsilon$ where ϵ is a small positive real (for example, 10^{-5}). Our new approximation becomes:

$$\tilde{\omega}_L^{DFT+DF}(i) = \eta^i \cdot \tilde{\omega}_L^{DFT}(i) = \frac{1}{N}\sum_{k=1}^{L} \psi(i_k)(\eta e^{\frac{2\pi j}{N} i_k})^i. \quad (6)$$

By incorporating this damping factor, we have that $\lim_{i\to+\infty} \tilde{\omega}_L^{DFT+DF}(i) = 0$. Especially, $\tilde{\omega}_L^{DFT+DF}(i) \leq \epsilon$ for $i > \alpha N$.

3. (Initial Scaling (IS)) Use of the damping factor gives a biased approximation when i is small (see Figure 3(i)). Taking the step function as an example, $\tilde{\omega}_L^{DFT+DF}(i)$ is approximately η^i for $0 \leq i < N$ instead of 1. To rectify this, we initially perform DFT on a different sequence $\hat{\omega}(i) = \eta^{-i}\omega(i)$ (rather than $\omega(i)$) on domain $\in [0, aN]$. This gives us an unbiased approximation, which we denote by $\tilde{\omega}^{DFT+DF+IS}$.

4. (Extending and Shifting (ES)) This trick is in particular tailored for optimizing the approximation performance for ranking functions. DFT does not perform well at discontinuous points, specifically at $i = 0$ (the left boundary), which can significantly affect the ranking approximation. To handle this, we extrapolate ω to make it continuous around 0. Let the resulting function be $\bar{\omega}$ which is defined on $[-bN, +\infty]$ for small ¿0. Again, taking the step function for example, $\bar{\omega}(i) = \begin{cases} 1, & -bN \leq i < N; \\ 0, & i \geq N. \end{cases}$ Then, we shift $\bar{\omega}(i)$ rightwards by bN to make its domain lie entirely in positive axis, do initial scaling and perform DFT on the resulting sequence. We denote the approximation of the resulting sequence by $\tilde{\omega}'(i)$(by performing (6)). For the approximation of original $\omega(i)$ values, we only need to do corresponding leftward shifting , namely $\tilde{\omega}^{DFT+DF+IS+ES}(i) = \tilde{\omega}'(i + bN)$. We can see from Figure 3(i) that DFT+DF+IS+ES produces a much better approximation than others around $i = 0$.

Figures 3(i) and (ii) illustrate the efficacy of our approximation technique for the step function. As we can see, we are able to approximate that function very well with just 20 or 30 coefficients. Figure 3(iii) and (iv) show the approximations for a piecewise linear function and an arbitrarily generated continuous function respectively, both of which are much easier to approximate than the step function.

5.2 Learning a PRF^ω or PRF^e Function

Next we address the question of how to learn the weights of a PRF^ω function or the α for a single PRF^e function from user preferences. To learn a linear combination of PRF^e functions, we first learn a PRF^ω function and then approximate it as above.

Prior work on learning ranking functions (e.g., [22, 28, 5, 10]) assumes that the user preferences are provided in the form of a set of pairs of tuples, and for each pair, we are told which tuple is ranked higher. Our problem differs slightly from this prior work in that, the features that we use to rank the tuples (i.e., $\mathsf{Pr}(r(t) = i), i = 1, \ldots, n$) cannot be computed for each tuple individually, but must be computed for the entire dataset (since the features for a tuple depend on the other tuples in the dataset). Hence, we assume that we are instead given a small sample of the tuples, and the user ranking for all those tuples. We compute the features assuming this sample constitutes the entire relation, and learn a ranking function accordingly, with the goal to find the parameters (the weights w_i for PRF^ω or the parameter α for PRF^e) that minimizes the number of disagreements with the provided ranking over the samples.

Given this, the problem of learning PRF^ω is identical to the problem addressed in the prior work, and we utilize the algorithm based on *support vector machines (SVM)* [28] in our experiments.

On the other hand, we are not aware of any work that has addressed learning a ranking function like PRF^e. We use a simple binary search-like heuristic to find the optimal real value of α that minimizes the Kendall distance between the user-specified ranking and the ranking according to $PRF^e(\alpha)$. In other words, we try to find $\arg\min_{\alpha\in[0,1]}(\mathsf{dis}(\sigma, \sigma(\alpha)))$ where $\mathsf{dis}()$ is the Kendall distance between two rankings, σ is the ranking for the given sample and $\sigma(\alpha)$ is the one obtained by using $PRF^e(\alpha)$ function. Suppose we want to find the optimal a within the interval $[L, U]$ now. We first compute $\mathsf{dis}(\sigma, \sigma(L + i \cdot \frac{U-L}{10})$ for $i = 1, \ldots, 9$ and find i for which the distance is the smallest. Then we reduce our search range to $[\max(L, L + (i-1) \cdot \frac{U-L}{10}, \min(U, L + (i+1) \cdot \frac{U-L}{10})]$ and repeat the above recursively. Although this algorithm can only converge to a local minimum, in our experimental study,

we observed that all of the prior ranking functions exhibit a uni-valley behavior (Section 7), and in such cases, this algorithm finds the global optimal.

6 PRF Computation for Arbitrary Correlations

Among many models for capturing the correlations in a probabilistic database, graphical models (Markov or Bayesian networks) perhaps represent the most systematic approach [33]. The appeal of graphical models stems both from the pictorial representation of the dependencies, and a rich literature on doing inference over them. In this section, we extend our generating function-based algorithm for computing PRF to handle correlations represented using a graphical model. The resulting algorithm is a non-trivial dynamic program over the junction tree of the graphical model, combined with the generating function method. Our main result is that we can compute the PRF function in polynomial time if the junction tree of the graphical model has bounded treewidth. It is worth noting that this result can not subsume our algorithm for and/xor trees (Section 4.2) since the treewidth of the moralized graph of a probabilistic and/xor tree may not be bounded.

Definitions: We start with briefly reviewing some notations and definitions related to graphical models and junction trees. Let $T = \{t_1, t_2, \ldots, t_n\}$ be the set of tuples in D_T, sorted in an non-increasing order of their score values. For each tuple t in T, we associate an indicator random variable X_t, which is 1, if t is present and 0 otherwise. Let $\mathcal{X} = \{X_{t_1}, \ldots, X_{t_n}\}$ and $\mathcal{X}_i = \{X_{t_1}, \ldots, X_{t_i}\}$. The correlations among these variables may be represented using either a directed or an undirected graphical model; we however assume that we are provided with an equivalent junction tree over the variables (which can be constructed using standard algorithms [15]).

Let $\mathcal{T}$ be a tree with each node v associated with a subset $C_v \subseteq \mathcal{X}$. We say $\mathcal{T}$ is a *junction tree* if any intersection $C_u \cap C_v$ for any $u, v \in \mathcal{T}$ is contained in C_w for every node w on the unique path between u and v in $\mathcal{T}$ (this is called the *running intersection property*). The treewidth of a junction tree is defined to be $\max_{v \in \mathcal{T}} |C_v| - 1$. Denote $S_{u,v} = C_v \cap C_u$ for each $(u, v) \in \mathcal{T}$. We call $S_{u,v}$ a *separator* since the removal of $S_{u,v}$ disconnects the graphical model.

Example 4 *Let $T = \{t_1, t_2, t_3, t_4, t_5\}$. Figure 4 (i) and (ii) show the (undirected) graphical model and the corresponding junction tree $\mathcal{T}$.$\mathcal{T}$ has four nodes v_1, v_2, v_3, v_4. $C_{v_1} = \{X_{t_4}, X_{t_5}\}$, $C_{v_2} = \{X_{t_4}, X_{t_3}\}$, $C_{v_2} = \{X_{t_3}, X_{t_1}\}$ and $C_{v_4} = \{X_{t_3}, X_{t_2}\}$. The treewidth of $\mathcal{T}$ is 1. We have, $S_{v_1,v_2} = \{X_4\}$, $S_{v_2,v_3} = \{X_3\}$ and $S_{v_2,v_4} = \{X_3\}$.*

We associate each clique C_v (and each separator $S_{u,v}$) with a *potential* $\pi_v(C_v)$ (resp. $\mu_{u,v}(S_{u,v})$), which is a function over all variables $X_{t_i} \in C_v$ ($X_{t_i} \in S_{u,v}$) and represents the correlations among those variables. Without loss of generality, we assume that the potentials are *calibrated*, that is, the potential corresponding to a clique (separator) is exactly the joint probability distribution over the variables in that clique (separator). Given a junction tree with arbitrary potentials, calibrated potentials can be computed using the message passing algorithm [15].

For a set of variables S, we use $\mathsf{Pr}(S)$ to denote the joint probability distribution over those variables. Then the joint probability distribution of $\mathcal{X}$, whose correlations can be captured using a calibrated junction tree $\mathcal{T}$, can be written as:

$$\mathsf{Pr}(\mathcal{X}) = \frac{\prod_{v \in \mathcal{T}} \pi_v(C_v)}{\prod_{(u,v) \in \mathcal{T}} \mu_{u,v}(S_{u,v})} = \frac{\prod_{v \in \mathcal{T}} \mathsf{Pr}(C_v)}{\prod_{(u,v) \in \mathcal{T}} \mathsf{Pr}(S_{u,v})}$$

Algorithm Sketch: Our dynamic programming-based algorithm computes $\mathsf{Pr}(r(t_i) = h)$ given any $t_i \in T$, for all $1 \leq h \leq n$, in polynomial time if the treewidth of $\mathcal{T}$ is bounded by a constant. To compute $\Upsilon(t_i)$ for all i, we need to run the algorithm n times, once for each t_i.

The algorithm begins by rooting $\mathcal{T}$ at a node root such that $X_{t_i} \in C_{\mathsf{root}}$. The dynamic program then runs bottom up, from the leaves to the root of the junction tree $\mathcal{T}$. Let $\mathcal{T}_v$ denote the subtree rooted at a node v in the junction tree. For each node v in $\mathcal{T}$, we maintain a polynomial number of states (which we call *configurations*). Each state represents a random event of the subtree $\mathcal{T}_v$. In each step, we compute the probability of one state of node v based on the probabilities of the states of v's children which we assume have been already computed. By aggregating a constant number of probabilities of the root's states, we can get $\mathsf{Pr}(r(t_i) = h)$.

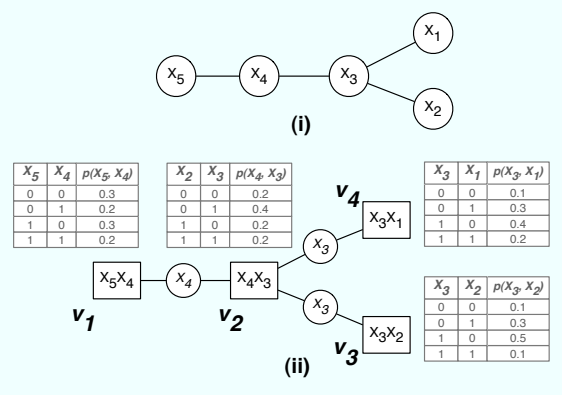

Figure 4: (i) A graphical model; (ii) A junction tree for the model along with the (calibrated) potentials.

Specifically, for each such node, we recursively compute the following quantities:

$$\mathsf{Pr}(\tilde{\mathcal{K}}_v^i) = \mathsf{Pr}((\tilde{\sigma}_v, \tilde{\theta}_v^i)), \qquad \forall \tilde{\sigma}_v \in \{0,1\}^{|C_v|}, \forall 0 \leq \tilde{\theta}_v^i \leq n$$

which is the probability that the variables in C_v take the values indicated by the boolean vector $\tilde{\sigma}_v$ (called a *configuration*) and the number of variables in $\mathcal{T}_v \cap \mathcal{X}_i$ is exactly equal to $\tilde{\theta}_v^i$.

After computing all of these values using dynamic programming, we can compute $\mathsf{Pr}(r(t_i) = h), \forall h$, as:

$$\mathsf{Pr}(r(t_i) = h) = \sum_{\substack{\tilde{\theta}^i_{\mathsf{root}} = h \\ \tilde{\sigma}_{\mathsf{root}}: \tilde{\sigma}_{\mathsf{root}}[t_i] = 1,}} \mathsf{Pr}(\tilde{\mathcal{K}}^i_{\mathsf{root}})$$

In other words, we compute the total probability that $X_{t_i} = 1$, and that exactly h variables in $\mathcal{X}_i$ are equal to 1 (i.e., exactly $h-1$ tuples ranked above t_i are present in the possible world).

Given the above framework, we can construct a recursive formula for computing $\mathsf{Pr}(\tilde{\mathcal{K}}_v^i)$ from (1) the computed values for the children of the node v, and (2) the joint probability distributions corresponding to the clique and its children. However it is not computationally feasible to evaluate that recursion formula directly. Instead we develop a generating functions-based algorithm for that purpose, which allows to efficiently compute $\mathsf{Pr}(\tilde{\mathcal{K}}_v^i)$ for all nodes v in the junction tree.

Example 5 *Let $T = \{t_1, t_2, t_3, t_4, t_5\}$. Figures 4 (i) and (ii) show an undirected graphical model over the corresponding variables and a calibrated junction tree $\mathcal{T}$ over them. Suppose we want to compute $\mathsf{Pr}(r(t_5) = 3)$. We root the junction tree at the clique v_1. Further, assume $\mathsf{score}(t_1) > \mathsf{score}(t_2) > \ldots > \mathsf{score}(t_5)$.*

Then, to compute $\mathsf{Pr}(r(t_5) = 3)$, the valid configurations for the node v_1 are $k_1 = (\langle 1, 0\rangle, 3)$ and $k_2 = (\langle 1, 1\rangle, 3)$ (the first index corresponds to X_{t_5}, the second to X_{t_4}), giving us:

$$\mathsf{Pr}(r(t_5) = 3) = \mathsf{Pr}\Big(\tilde{\mathcal{K}}^5_{v_1} = (\langle 1, 1\rangle, 3)\Big) + \mathsf{Pr}\Big(\tilde{\mathcal{K}}^5_{v_1} = (\langle 1, 0\rangle, 3)\Big)$$

The first term on the r.h.s. corresponds to the probability that both X_{t_5} and X_{t_4} are set to 1, and exactly one other term in $X_{t_1}, \ldots, X_{t_5}$ is set to 1. Whereas the second term corresponds to $X_{t_5} = 1, X_{t_4} = 0$, and exactly two of $X_{t_3}, X_{t_2}, X_{t_1}$ set to 1. □

Algorithm Details: We begin with some notations first. Let $\mathcal{T}_v$ be the subtree rooted at the vertex $v \in \mathcal{T}$. For a subset $S \subseteq \mathcal{X}$, let $\tilde{\sigma}_S \in \{0,1\}^{|S|}$ be a boolean vector of size $|S|$ which we call *the existential configuration* of S. Essentially, $\tilde{\sigma}_S$ represent an outcome of the random variables in S. For ease of notation, we write $\tilde{\sigma}_v, \tilde{\sigma}_{u,v}$ and $\tilde{\sigma}_{\mathcal{T}_v}$ in place of $\tilde{\sigma}_{C_v}, \tilde{\sigma}_{S_{u,v}}$ and $\tilde{\sigma}_{\cup_{u \in \mathcal{T}_v} C_u}$, respectively. Let $|\tilde{\sigma}|$ be the number of 1's in vector $\tilde{\sigma}$. We define $\tilde{\sigma}^i_S$ to be the restriction of $\tilde{\sigma}_S$ which only contains entries in $S \cap \mathcal{X}_i$. We say two existential configurations $\tilde{\sigma}_{S_1}$ and $\tilde{\sigma}_{S_2}$ are *consistent* with each other if they agree with each other on their intersection $S_1 \cap S_2$, denoted by $\tilde{\sigma}_{S_1} \sim \tilde{\sigma}_{S_2}$.

Suppose $v_1, v_2, \ldots, v_l$ be the children of v in $\mathcal{T}$. Before providing the recursion formula, we need the following simple fact, which follows from the running intersection property of junction tree.

$\mathcal{X}_i$	$\{X_{t_1}, \ldots, X_{t_i}\}$
C_v	$\subseteq \mathcal{X}$, for each node $v \in \mathcal{T}$
$S_{u,v}$	$C_u \cap C_v$, for each edge $(u, v) \in \mathcal{T}$
$\mathcal{T}_v$	Subtree rooted at $v \in \mathcal{T}$
$\tilde{\sigma}_S$	$\{0,1\}^{\lvert S\rvert}$ for $S \subseteq \mathcal{X}$.
	$\tilde{\sigma}_S[t_i] = 1$ indicates $X_{t_i} = 1$ for $X_{t_i} \in S$
$\tilde{\sigma}^i_S$	restriction of $\tilde{\sigma}_S$ on entries in $S \cap \mathcal{X}_i$
$\tilde{\theta}_v$	# nodes in $\mathcal{T}_v$
$\tilde{\mathcal{K}}_v$	$(\tilde{\sigma}_v, \tilde{\theta}_v)$

Table 2: Some Useful Notations

Lemma 1 *1. If $\tilde{\sigma}_{v_j} \sim \tilde{\sigma}_v$ and $\tilde{\sigma}_{v_k} \sim \tilde{\sigma}_v$ in $\mathcal{T}$, then $\tilde{\sigma}_{v_j} \sim \tilde{\sigma}_{v_k}$.*

2. If $\tilde{\sigma}_{v_j} \sim \tilde{\sigma}_v$, then $\tilde{\sigma}_{T_{v_j}} \sim \tilde{\sigma}_v$.

Proof: Both parts follow directly from the running intersection property of junction trees. More specifically, the first part can be seen from the fact that $C_{v_i} \cap C_{v_j} \subseteq C_v$ and the second holds because $C_u \cap C_v \subseteq C_{v_i}$ for any $u \in T_{v_i}$. □

We use shorthand notation $[\tilde{\sigma}_{v_j}]$ to denote the set $\tilde{\sigma}_{v_1}, \ldots, \tilde{\sigma}_{v_l}$, $[\tilde{\theta}_{v_j}]$ to denote $\tilde{\theta}_{v_1}, \ldots, \tilde{\theta}_{v_l}$ and $[\tilde{\mathcal{K}}_{v_j}]$ to denote $\tilde{\mathcal{K}}_{v_1}, \ldots, \tilde{\mathcal{K}}_{v_l}$. We say $\tilde{\mathcal{K}}_v$ is consistent with $[\tilde{\mathcal{K}}_{v_j}]$, denoted by $\tilde{\mathcal{K}}_v \sim [\tilde{\mathcal{K}}_{v_j}]$, if there exists a configuration $\tilde{\sigma}_{\mathcal{T}_v}$ such that $\tilde{\sigma}_{T_v} \sim \tilde{\sigma}_v \sim [\tilde{\sigma}_{v_j}]$, $|\tilde{\sigma}^i_{\mathcal{T}_v}| = \tilde{\theta}_v$ and its restriction on T_{v_j}, denoted by $\tilde{\sigma}_{\mathcal{T}_{v_j}}$, satisfies $|\tilde{\sigma}^i_{T_{v_j}}| = \tilde{\theta}_{v_j}$ for all j. The following lemma characterizes the condition when such consistent configurations can be obtained. We use this lemma for the construction of generating functions.

Lemma 2 *$[\tilde{\mathcal{K}}_{v_j}] \sim \tilde{\mathcal{K}}_v$ if and only if $[\tilde{\sigma}_{v_j}] \sim \tilde{\sigma}_v$ and $|\tilde{\sigma}^i_v| - \sum_{j=1}^l |\tilde{\sigma}^i_{v,v_j}| = \tilde{\theta}_v - \sum_{j=1}^l \tilde{\theta}_{v_j}$*

Proof: Note that $\tilde{\theta}_v = |(\cup_{j=1}^l \tilde{\sigma}^i_{v_j} \cup \tilde{\sigma}^i_v)|$ and $\sum_{j=1}^l \tilde{\theta}_{v_j} = \sum_{j=1}^l |\tilde{\sigma}^i_{v_j}|$. Thus we only need to prove, $|(\cup_{j=1}^h \tilde{\sigma}^i_{v_j} \cup \tilde{\sigma}^i_v)| = \sum_{j=1}^h |\tilde{\sigma}^i_{v_j}| + |\tilde{\sigma}^i_v| - \sum_{j=1}^h |\tilde{\sigma}^i_{v_j} \cap \tilde{\sigma}^i_v|$, for $h = l$. We prove this by induction on h. For $h = 0$, it holds trivially. Let us assume, it holds for values up to h. We have,

$$\begin{aligned}
|(\cup_{j=1}^{h+1} \tilde{\sigma}^i_{v_j} \cup \tilde{\sigma}^i_v)| &= |(\cup_{j=1}^{h} \tilde{\sigma}^i_{v_j} \cup \tilde{\sigma}^i_v)| + |\tilde{\sigma}^i_{v_{h+1}}| - |((\cup_{j=1}^{h} \tilde{\sigma}^i_{v_j} \cup \tilde{\sigma}^i_v) \cap \tilde{\sigma}^i_{v_{h+1}})| \\
&= \sum_{j=1}^{h} |\tilde{\sigma}^i_{v_j}| + |\tilde{\sigma}^i_v| - \sum_{j=1}^{h} |\tilde{\sigma}^i_{v_j} \cap \tilde{\sigma}^i_v| + |\tilde{\sigma}^i_{v_{h+1}}| - |\tilde{\sigma}^i_{v_{h+1}} \cap \tilde{\sigma}^i_v| \\
&= \sum_{j=1}^{h+1} |\tilde{\sigma}^i_{v_j}| + |\tilde{\sigma}^i_v| - \sum_{j=1}^{h+1} |\tilde{\sigma}^i_{v_j} \cap \tilde{\sigma}^i_v|
\end{aligned}$$

The first equality follows from inclusion-exclusion principle and the second from the induction hypothesis and the running intersection property. So, the claim holds. □

The recursion formula of dynamic programming is given by

$$\mathsf{Pr}(\tilde{\mathcal{K}}_v) \quad = \quad \frac{\pi(\tilde{\sigma}_v)}{\prod_{j=1}^{l} \mu_{v,v_j}(\tilde{\sigma}_{v,v_j})} \cdot \sum_{[\tilde{\mathcal{K}}_{v_j}] \sim \tilde{\mathcal{K}}_v} \prod_{j=1}^{l} \mathsf{Pr}(\mathcal{K}_{v_j}) \tag{7}$$

The following theorem establishes the correctness of the computed recursion for dynamic programming.

Theorem 2 *The recursion formula (7) correctly computes the probability*

$$\mathsf{Pr}(\tilde{\mathcal{K}}_v) = \sum_{\tilde{\sigma}_{\mathcal{T}_v} : |\tilde{\sigma}^i_{\mathcal{T}_v}| = \tilde{\theta}_v} \mathsf{Pr}(\tilde{\sigma}_{\mathcal{T}_v}).$$

Proof: We prove this theorem by induction on the tree structure. Assume it holds for all v's children, i.e,

$$\mathsf{Pr}(\tilde{\mathcal{K}}_{v_j}) = \sum_{\tilde{\sigma}_{T_{v_j}}:|\tilde{\sigma}^i_{T_{v_j}}|=\tilde{\theta}_{v_j}} \mathsf{Pr}(\tilde{\sigma}_{T_{v_j}}).$$

for all j,$\sigma_{T_{v_j}}$ and θ_{v_j}. Then, plugging them into the recursion, we get

$$\begin{aligned}
\mathsf{Pr}(\tilde{\mathcal{K}}_v) &= \frac{\pi(\tilde{\sigma}_v)}{\prod_{j=1}^{l} \mu_{v,v_j}(\tilde{\sigma}_{v,v_j})} \cdot \sum_{[\tilde{\mathcal{K}}_{v_j}]\sim\tilde{\mathcal{K}}_v} \prod_{j=1}^{l} \mathsf{Pr}(\tilde{\mathcal{K}}_{v_j}) \\
&= \frac{\pi(\tilde{\sigma}_v)}{\prod_{j=1}^{l} \mu_{v,v_j}(\tilde{\sigma}_{v,v_j})} \sum_{[\tilde{\mathcal{K}}_{v_j}]\sim\tilde{\mathcal{K}}_v} \left(\prod_{j=1}^{l} \sum_{\tilde{\sigma}_{T_{v_j}}:|\tilde{\sigma}^i_{T_{v_j}}|=\tilde{\theta}_{v_j}} \mathsf{Pr}(\tilde{\sigma}_{T_{v_j}}) \right) \\
&= \sum_{[\tilde{\mathcal{K}}_{v_j}]\sim\tilde{\mathcal{K}}_v} \left(\sum_{\tilde{\sigma}_{T_{v_j}}\sim\tilde{\sigma}_{T_v}:|\tilde{\sigma}^i_{T_{v_j}}|=\tilde{\theta}_{v_j}} \left(\frac{\pi(\tilde{\sigma}_v) \prod_{j=1}^{l} \mathsf{Pr}(\tilde{\sigma}_{T_{v_j}})}{\prod_{j=1}^{l} \mu_{v,v_j}(\tilde{\sigma}_{v,v_j})} \right) \right) \\
&= \sum_{[\tilde{\mathcal{K}}_{v_j}]\sim\tilde{\mathcal{K}}_v} \left(\sum_{\tilde{\sigma}_{T_{v_j}}\sim\tilde{\sigma}_{T_v}:|\tilde{\sigma}^i_{T_{v_j}}|=\tilde{\theta}_{v_j}} \mathsf{Pr}(\tilde{\sigma}_{T_v}) \right), \text{ see explanation below} \\
&= \sum_{\tilde{\sigma}_{T_v}:|\tilde{\sigma}^i_{T_v}|=\tilde{\theta}_v} \mathsf{Pr}(\tilde{\sigma}_{T_v})
\end{aligned}$$

The last equality holds because of the definition of $[\tilde{\mathcal{K}}_{v_j}] \sim \tilde{\mathcal{K}}_v$. Now, we prove the fourth equality $\pi(\tilde{\sigma}_v) \prod_{j=1}^{l} \frac{P(\tilde{\sigma}_{T_{v_j}})}{\mu_{v,v_j}(\sigma_{v,v_j})} = \mathsf{Pr}(\tilde{\sigma}_{T_v})$ for $\tilde{\sigma}_v \sim [\tilde{\sigma}_{v_j}]$. Since C_v is a separator of different T_{v_j}s, so by local Markov assumptions ($T_{v_1} \perp T_{v_j} \perp \ldots | C_i$), i.e, $T_{v_1}, T_{v_2}, \ldots$ are conditionally independent of each other given C_i. Written in probability form, it is $\mathsf{Pr}(\tilde{\sigma}_{T_v}|\tilde{\sigma}_v) = \prod_{j=1}^{l} \mathsf{Pr}(\tilde{\sigma}_{T_{v_j}}|\tilde{\sigma}_v)$. Therefore, for any $\tilde{\sigma}_{T_v}$ such that $\tilde{\sigma}_{T_v} \sim [\tilde{\sigma}_{T_{v_j}}]$, we have

$$\mathsf{Pr}(\tilde{\sigma}_{T_v}) = \mathsf{Pr}(\tilde{\sigma}_{T_v}|\tilde{\sigma}_v)\mathsf{Pr}(\tilde{\sigma}_v) = \pi(\tilde{\sigma}_v) \prod_{j=1}^{l} \mathsf{Pr}(\tilde{\sigma}_{T_{v_j}}|\tilde{\sigma}_v) = \pi(\tilde{\sigma}_v) \prod_{j=1}^{l} \frac{\mathsf{Pr}(\tilde{\sigma}_{T_{v_j}})}{\mathsf{Pr}(\tilde{\sigma}_{v,v_j})} = \pi(\tilde{\sigma}_v) \prod_{j=1}^{l} \frac{\mathsf{Pr}(\tilde{\sigma}_{T_{v_j}})}{\mu_{v,v_j}(\tilde{\sigma}_{v,v_j})}$$

□

Example 6 *Suppose we are considering node v_2 and want to compute $\mathsf{Pr}\Big(\tilde{\mathcal{K}}_{v_2}(\langle 1,0\rangle, 2)\Big)$. Using Lemma 2, we know consistent configurations for node v_3 and v_4 are $(\langle 0,0\rangle, 0), (\langle 0,1\rangle, 1)$ and $(\langle 0,1\rangle, 1), (\langle 0,0\rangle, 0)$. Thus,*

$$\begin{aligned}
&\mathsf{Pr}(\tilde{\mathcal{K}}_{v_2}(\langle 1,0\rangle, 2) \\
=\ & \frac{\mathsf{Pr}(X_4 = 1 \wedge X_3 = 0)}{\mathsf{Pr}(X_3 = 0)^2} \Big(\mathsf{Pr}(\tilde{\mathcal{K}}_{v_3}(\langle 0,0\rangle, 0)\mathsf{Pr}(\tilde{\mathcal{K}}_{v_4}(\langle 0,1\rangle, 1) \\
& + \mathsf{Pr}(\tilde{\mathcal{K}}_{v_3}(\langle 0,1\rangle, 1)\mathsf{Pr}(\tilde{\mathcal{K}}_{v_4}(\langle 0,0\rangle, 0) \Big) \\
=\ & \frac{0.2}{0.4^2}(0.1 \times 0.3 + 0.1 \times 0.3) = 0.075
\end{aligned}$$

□

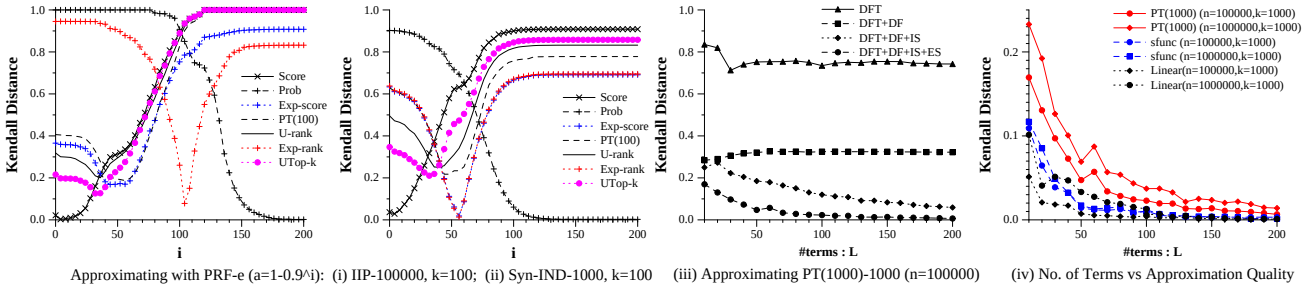

Approximating with PRF-e (a=1-0.9^i): (i) IIP-100000, k=100; (ii) Syn-IND-1000, k=100 (iii) Approximating PT(1000)-1000 (n=100000) (iv) No. of Terms vs Approximation Quality

Figure 5: (i, ii) Comparing PRF^e with other ranking functions for varying values of α; (iii) Approximating $PT(1000)$ using a linear combination of PRF^e functions; (iv) Approximation quality for three ranking functions for varying number of exponentials.

However, the number of different possible $[\tilde{\mathcal{K}}_{v_j}]$ consistent with $[\tilde{\mathcal{K}}_v]$ may be exponential. So it is still not clear how (7) can be evaluated in polynomial time. Suppose we have already computed $\mathsf{Pr}(\tilde{\mathcal{K}}_{v_j})$ for all v_j and all different possible configurations of $\tilde{\mathcal{K}}_{v_j}$ and now we want to compute $\mathsf{Pr}(\tilde{\mathcal{K}}_v)$ for a fixed $\tilde{\mathcal{K}}_v = (\tilde{\sigma}_v, \tilde{\theta}_v)$. Note that when $\tilde{\mathcal{K}}_v = (\tilde{\sigma}_v, \tilde{\theta}_v)$ is fixed, $\tilde{\sigma}_{v,v_j}$ becomes fixed $\forall j = 1, 2, .., l$, since $\tilde{\sigma}_{v,v_j}$ contains the configurations of only a subset of bits in $\tilde{\sigma}_v$. Thus, $I = \tilde{\theta}_v - |\tilde{\sigma}^i_v| + \sum_{j=1}^{l} |\tilde{\sigma}^i_{v,v_j}|$ (see Lemma 2) does not depend on the actual values of $[\tilde{\sigma}_{v_j}]$ if $[\tilde{\sigma}_{v_j}] \sim \tilde{\sigma}_v$. So we can treat I as a constant. Similarly, the term $\frac{\pi(\tilde{\sigma}_v)}{\prod_{j=1}^{l} \mu_{v,v_j}(\tilde{\sigma}_{v,v_j})}$ is also fixed.

Thus we can rewrite,

$$\sum_{[\tilde{\mathcal{K}}_{v_j}] \sim \tilde{\mathcal{K}}_v} \prod_{j=1}^{l} \mathsf{Pr}(\tilde{\mathcal{K}}_{v_j}) = \sum_{[\tilde{\sigma}_{v_j}] \sim \tilde{\sigma}_v} \sum_{\sum_j \tilde{\theta}_{v_j} = I} \prod_{j=1}^{l} \mathsf{Pr}(\tilde{\mathcal{K}}_{v_j}) = \sum_{\sum_j \tilde{\theta}_{v_j} = I} \prod_{j=1}^{l} \left(\sum_{[\tilde{\sigma}_{v_j}] \sim \tilde{\sigma}_v} \mathsf{Pr}(\tilde{\mathcal{K}}_{v_j}) \right) = \sum_{\sum_j \tilde{\theta}_{v_j} = I} \prod_{j=1}^{l} a_{v_j, \tilde{\theta}_{v_j}}$$

where $a_{v_j,\tilde{\theta}_{v_j}} = \sum_{\tilde{\sigma}_{v_j} \sim \tilde{\sigma}_v} \mathsf{Pr}(\tilde{\mathcal{K}}_{v_j})$. We know $\forall j = 1, \ldots, l \; |\tilde{\sigma}_{v_j}| \leq tw$. So if the treewidth tw is a constant, there are at most 2^{tw} (also a constant) different $\tilde{\sigma}_{v_j}$s. Thus, each $a_{v_j,\tilde{\theta}_{v_j}}$ can be computed in constant time for all $1 \leq j \leq l$ and $0 \leq \tilde{\theta}_{v_j} \leq n$. Consider the following generating function:

$$\mathcal{F}(x) = \prod_{j=1}^{l} \left(\sum_{k=0}^{n} a_{v_j,k} x^k \right)$$

It is not hard to see that the coefficient of $x^{\tilde{\theta}_v}$ is exactly equal to $\sum_{[\tilde{\mathcal{K}}_{v_j}] \sim \tilde{\mathcal{K}}_v} \prod_{j=1}^{l} \mathsf{Pr}(\tilde{\mathcal{K}}_{v_j})$. Therefore, we only have to evaluate $a_{v_j,k}$ for all v_j and k, and then construct and expand $\mathcal{F}(x)$ appropriately.

Running Time: We need to run our dynamic program n times for each tuple t_i. In the dynamic program, we need to construct and expand $O(2^{|C_v|})$ generating functions for each node v in $\mathcal{T}$, since one generating function is able to generate $\mathsf{Pr}\Big((\tilde{\sigma}_v, \tilde{\theta}_v)\Big)$ for a fixed $\tilde{\sigma}_v$ and all $1 \leq \tilde{\theta}_v \leq n$ and we have $2^{|C_v|}$ different $\tilde{\sigma}_v$ values. To construct and expand a generating function at node v takes $O(\mathsf{deg}(v)(n2^{|C_{v_j}|} + n^2))$ time. Specifically, computing each $a_{v_j,i}$ takes $O(2^{|C_{v_j}|})$ and expanding the polynomial takes $O(\mathsf{deg}(v)n^2)$ time. Therefore, the time complexity is $O(2^{tw}(n2^{tw} + n^2)|\mathcal{T}|)$ for each execution of dynamic program, resulting in an overall time complexity of $O(2^{tw}n^2(2^{tw} + n)|\mathcal{T}|)$, where tw is the treewidth of the junction tree $\mathcal{T}$.

7 Experimental Study

We conducted an extensive empirical study over several real and synthetic datasets to illustrate: (a) the diverse and conflicting behavior of different ranking functions proposed in the prior literature, (b) the effectiveness of our param-

eterized ranking functions, especially PRF^e, at approximating other ranking functions, and (c) the scalability of our new generating functions-based algorithms for exact and approximate ranking. We discussed the results supporting (a) in Section 3.2. In this section, we focus on (b) and (c).

Datasets: We mainly use the International Ice Patrol (IIP) Iceberg Sighting Dataset[2] for our experiments. This dataset was also used in prior work on ranking in probabilistic databases [27, 23]. The database contains a set of *iceberg sighting records*, each of which contains the location (*latitude, longitude*) of the iceberg, and the *number of days* the iceberg has drifted, among other attributes. Detecting the icebergs that have been drifting for long periods is crucial, and hence we use the number of days drifted as the ranking score. The sighting record is also associated with a *confidence-level* attribute according to the source of sighting: R/V (radar and visual), VIS (visual only), RAD (radar only), SAT-LOW (low earth orbit satellite), SAT-MED (medium earth orbit satellite), SAT-HIGH (high earth orbit satellite), and EST (estimated). We converted these six confidence levels into probabilities 0.8, 0.7, 0.6, 0.5, 0.4, 0.3, and 0.4 respectively. We added a very small Gaussian noise to each probability so that ties could be broken. There are nearly a million records available from 1960 to 2007; we created 10 different datasets for our experimental study containing $100,000$ (IIP-100,000) to $1,000,000$ (IIP-1,000,000) records, by uniformly sampling from the original dataset.

Along with the real datasets, we also use several synthetic datasets with varying degrees of correlations, where the correlations are captured using probabilistic and/xor trees. The tuple scores (for ranking) were chosen uniformly at random from $[0, 10000]$. The corresponding and/xor trees were also generated randomly starting with the root, by controlling the *height (L)*, the *maximum degree of the non-root nodes (d)*, and the *proportion of* $\veebar$ *and* $\wedge$ *nodes (X/A)* in the tree. Specifically, we use five such datasets:

(1) Syn-IND (independent tuples), (2) Syn-XOR (L=2,X/A=∞,d=5), (3) Syn-LOW (L=3,X/A=10,d=2), (4) Syn-MED (L=5,X/A=3,d=5), and (5) Syn-HIGH (L=5,X/A=1,d=10).

For Syn-IND, the tuple existence probabilities were chosen uniformly at random from $[0, 1]$. Note that the Syn-XOR dataset, with height set to 2 and no $\wedge$ nodes, exhibits only mutual exclusivity correlations (mimicking the x-tuples model [32, 36]), whereas the latter three datasets exhibit increasingly more complex correlations.

Setup: We use the normalized Kendall distance (Section 3.2) for comparing two top-k rankings. All the algorithms were implemented in C++, and the experiments were run on a 2.4GHz Linux PC with 2GB memory.

7.1 Approximability of Ranking Functions

We begin with a set of experiments illustrating the effectiveness of our parameterized ranking functions at approximating other ranking functions. Due to space constraints, we focus on PRF^e here because it is significantly faster to rank according to a PRF^e function (or a linear combination of several PRF^e functions) than it is to rank according a PRF^ω function.

Figures 5 (i) and (ii) show the Kendall distance between the Top-100 answers computed using a specific ranking function and PRF^e for varying values of α, for the IIP-100,000 and Syn-IND-1000 datasets. For better visualization, we plot i on the x-axis, where $\alpha = 1 - 0.9^i$. The reason behind this is that the behavior of the PRF^e function changes rather drastically, and spans a spectrum of rankings, when α approaches 1. First, as we can see, the PRF^e ranking is close to ranking by *Score* alone for small values of α, whereas it is close to the ranking by *Probability* when α is close to 1 (in fact, for $\alpha = 1$, the PRF^e ranking is equivalent to the ranking of tuples by their existence probabilities)[3]. Second, we see that, for all other functions (*E-Score*, $PT(h)$, *U-Rank*, *Exp-Rank*), there exists a value of α for which the distance of that function to PRF^e is very small, indicating that PRF^e can indeed approximate those functions quite well. Moreover we observe that this "uni-valley" behavior of the curves justifies the binary search algorithm we advocate for learning the value of α in Section 5.2. Our experiments with other synthetic and real datasets indicated a very similar behavior by the ranking functions.

Next we evaluate the effectiveness of our approximation technique presented in Section 5.1. In Figure 5 (iii), we show the Kendall distance between the top-k answers obtained using $PT(h)$ (for $h = 1000, \mathrm{k} = 1000$) and using a linear combination of PRF^e functions found by our algorithms. As expected, the approximation using the vanilla

[2] http://nsidc.org/data/g00807.html

[3] On the other hand, for $\alpha = 0$, PRF^e ranks the tuples by their probabilities to be the Top-1 answer.

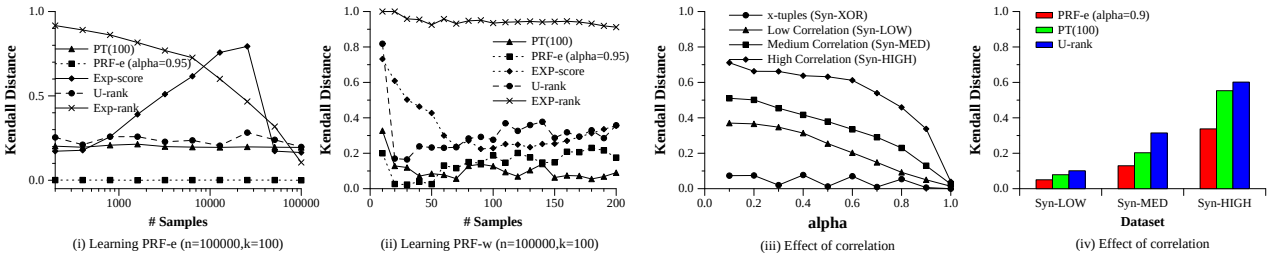

Figure 6: (i) Learning PRF^e from user preferences; (ii) Learning PRF^ω from user preferences; (iii) Effect of correlations on PRF^e ranking as a varies; (iv) Effect of correlations on PRF^e, *U-Rank* and $PT(h)$.

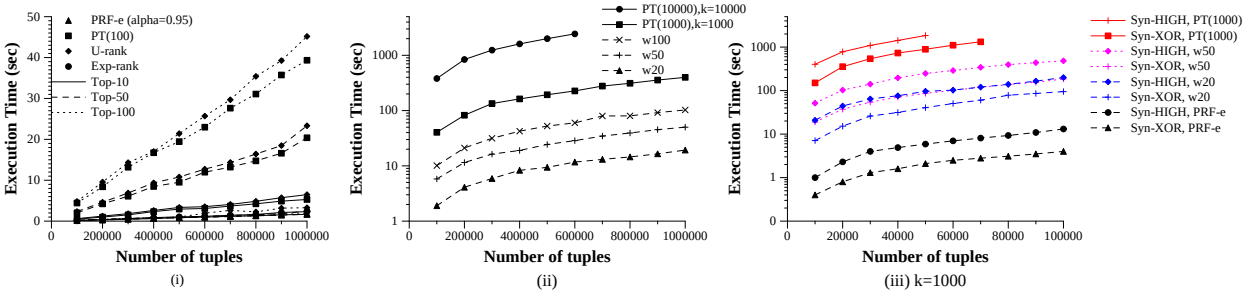

Figure 7: Experiments comparing the execution times of the ranking algorithms (note that the y-axis is log-scale for (ii) and (iii))

DFT technique is very bad, with the Kendall distance close to 0.8 indicating little similarity between the top-k answers. However, the approximation obtained using our proposed algorithm (indicated by DFT+DF+IS+ES curve) achieves a Kendall distance of less than 0.1 with just $L = 20$ exponentials.

In Figure 5 (iv), we compare the approximation quality (found by our algorithm DFT+DF+IS+ES) for three ranking functions for two datasets: IIP-100,000 with $\mathrm{k} = 100$, and IIP-1,000,000 dataset with $\mathrm{k} = 10000$. The ranking functions we compared were: (1) $PT(h)$ ($h = 1000$), (2) an arbitrary smooth function, $sfunc$, (see Figure 3(iv)), and (3) a linear function (Figure 3(iii)). We see that $L = 40$ suffices to bring the Kendall distance to < 0.1 in all cases. We also observe that smooth functions (for which the absolute value of the first derivative of the underlying continuous function is bounded by a small value) are usually easier to approximate. We only need $L = 20$ exponentials to achieve a Kendall distance less than 0.05 for $sfunc$. The Linear function is even easier to approximate. We also tested a few other continuous functions such as piecewise linear function and $f(x) = 1/x$, and found similar behavior. We omit those curves for clarity.

7.2 Learning Ranking Functions

Next we consider the issue of learning ranking functions from user preferences. Lacking real user preference data, we instead assume that the user ranking function, denoted *user-func*, is identical to one of: *E-Score*, $PT(h)$, *U-Rank*, *Exp-Rank*, or $PRF^e(\alpha = 0.95)$. We generate a set of user preferences by ranking a random sample of the dataset using *user-func* (thus generating five sets of user preferences). These are then fed to the learning algorithm, and finally we compare the Kendall distance between the learned ranking and the true ranking for the entire dataset.

In Figure 6(i), we plot the results for learning a single PRF^e function (i.e., for learning the value of α) using the binary search-like algorithm presented in Section 5.2. The experiment reveals that when the underlying ranking is done by PRF^e, the value of α can be learned perfectly. When one of $PT(h)$ or *U-Rank* is the underlying ranking function, the correct value a can be learned with a fairly small sample size, and increasing the number of samples does not help in finding a better α. On the other hand, *Exp-Rank* cannot be learned well by PRF^e unless the sample size approaches the total size of whole dataset. This phenomenon can be partly explained using Figure 5(i) in which the curves for $PT(h)$ and *U-Top* have a fairly smooth valley, while the one for *Exp-Rank* is very sharp and the region of α

values where the distance is low is extremely small ($[1-0.9^{90}, 1-0.9^{110}]$). Hence, the minimum point for *Exp-Rank* is harder to reach. Another reason is that *Exp-Rank* is quite sensitive to the size of the dataset, which makes it hard to learn it using a smaller-sized sample dataset. We also observe that while extremely large samples are able to learn *E-Score* well, the behavior of *E-Score* is quite unstable when the sample size is smaller.

Note that if we already know the form of the ranking function, we don't need to learn it in this fashion; we can instead directly find an approximation for it using our DFT-based algorithm.

In Figure 6 (ii), we show the results of an experiment where we tried to learn a PRF^ω function (using the SVM-lite package [28]). We keep our sample size ≤ 200 since SVM-lite becomes drastically slow with larger sample sizes. First we observe that $PT(h)$ and PRF^e can be learned very well from a small size sample (distance < 0.2 in most cases) and increasing the sample size does not benefit significantly. *U-Rank* can also be learned, but the approximation isn't nearly as good. This is because *U-Rank* can not be written as a single PRF^ω function. We observed similar behavior in our experiments with other datasets. Due to space constraints, we omit a further discussion on learning a PRF^ω function; the issues in learning such weighted functions have been investigated in prior literature, and if the true ranking function can be written as a PRF^ω function, then the above algorithm is expected to learn it well given a reasonable number of samples.

7.3 Effect of Correlations

Next we evaluate the behavior of ranking functions over probabilistic datasets modeled using probabilistic and/xor trees. We use the four synthetic correlated datasets, Syn-XOR, Syn-LOW, Syn-MED, and Syn-HIGH, for these experiments. For each dataset and each ranking function considered, we compute the rankings by considering the correlations, and by ignoring the correlations, and then compute the Kendall distance between these two (e.g., for PRF^e, we compute the rankings using **PROB-ANDOR-PRF-RANK** and **IND-PRF-RANK** algorithms). Figure 6(iii) shows the results for the PRF^e ranking function for varying α, whereas in Figure 6(iv), we plot the results for $PRF^e(\alpha = 0.9)$, $PT(100)$, and *U-Rank*. As we can see, on highly correlated datasets, ignoring the correlations can result in significantly inaccurate top-k answers. This is not as pronounced for the Syn-XOR dataset. This is because, in any group of tuples that are mutually exclusive, there are typically only a few tuples that may have sufficiently high probabilities to be part of the top-k answer; the rest of the tuples may be ignored for ranking purposes. Because of this, assuming tuples to be independent of each other does not result in significant errors. As α approaches 1, PRF^e tends to sort the tuples by probabilities, so all four curves in Figure 6(iii) become close to 0. We note that ranking by *E-Score* is invariant to the correlations, which is a significant drawback of that function.

7.4 Execution Times

Figure 7(i) shows the execution times for four ranking functions: PRF^e, $PT(h)$, *U-Rank* and *Exp-Rank*, for the IIP-datasets, for different dataset sizes and k. We note that the running time for PRF^ω is similar to that of $PT(h)$. As expected, ranking by PRF^e or *Exp-Rank* is very efficient (1000000 tuples can be ranked within 1 or 2 seconds). Indeed, after sorting the dataset in an non-decreasing score order, PRF^e needs only a single scan of the dataset, and *Exp-Rank* needs to scan the dataset twice. Execution times for $PT(h)$ and *U-Rank*-k increase linearly with h and k respectively and the algorithms become very slow for high h and k. The running times of both PRF^e and *Exp-Rank* are not significantly affected by k.

Figure 7(ii) compares the execution time for $PT(h)$ and its approximation using a linear combination of PRF^e functions (see Figure 5(iii)), for two different values of k. $w50$ indicates that 50 exponentials were used in the approximation (note that the approximate ranking, based on PRF^e, is insensitive to the value of k). As we can see, for large datasets and for higher values of k, exact computation takes several orders of magnitude more time to compute than the approximation. For example, the exact algorithm takes nearly 1 hour for $n = 500,000$ and $h = 10,000$ while the approximate answer obtained using $L = 50$ PRF^e functions takes only 24 seconds and achieves a Kendall distance 0.09.

For correlated datasets, the effect is even more pronounced. In Figure 7(iii), we plot the results of a similar experiment, but using two correlated datasets: Syn-XOR and Syn-HIGH. Note that the number of tuples in these datasets is smaller by a factor of 10. As we can see, our generating functions-based algorithms for computing PRF^e are highly efficient, even for datasets with high degrees of correlation. As above, approximation of the $PT(h)$ ranking

function using a linear combination of PRF^e functions is significantly cheaper to compute than using the exact algorithm.

Combined with the previous results illustrating that a linear combination of PRF^e functions can approximate other ranking functions very well, this validates the unified ranking approach that we propose in this paper.

8 Conclusions

In this paper we presented a unified framework for ranking over probabilistic databases, and presented several novel and highly efficient algorithms for answering top-k queries. Considering the complex interplay between probabilities and scores, instead of proposing a specific ranking function, we propose using two parameterized ranking functions, called PRF^ω and PRF^e, which allow the user to control the tuples that appear in the top-k answers. We developed novel algorithms for evaluating these ranking functions over large, possibly correlated, probabilistic datasets. We also developed an approach for approximating a ranking function using a linear combination of PRF^e functions thus enabling highly efficient, albeit approximate computation, and also for learning a ranking function from user preferences. Our work opens up many avenues for further research that we are planning to pursue. For instance, there may be other non-trivial subclasses of PRF functions, aside from PRF^e, that can be computed very efficiently. Understanding the behavior of various ranking functions and their relationships across probabilistic databases with diverse uncertainties and correlation structures also remains an important open problem in this area.

References

[1] E. Adar and C. Re. Managing uncertainty in social networks. *IEEE Data Eng. Bull.*, 2007.

[2] Periklis Andritsos, Ariel Fuxman, and Renee J. Miller. Clean answers over dirty databases. In *ICDE*, 2006.

[3] L. Antova, C. Koch, and D. Olteanu. From complete to incomplete information and back. In *SIGMOD*, 2007.

[4] G. Beylkin and L. Monzon. On approximation of functions by exponential sums. *Applied and Computational Harmonic Analysis*, 19:17–48, 2005.

[5] Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. Learning to rank using gradient descent. In *ICML*, pages 89–96, 2005.

[6] R. Cheng, D. Kalashnikov, and S. Prabhakar. Evaluating probabilistic queries over imprecise data. In *SIGMOD*, 2003.

[7] G. Cormode, F. Li, and K. Yi. Semantics of ranking queries for probabilistic data and expected ranks. In *ICDE*, 2009.

[8] Nilesh Dalvi and Dan Suciu. Efficient query evaluation on probabilistic databases. In *VLDB*, 2004.

[9] Nilesh Dalvi and Dan Suciu. Management of probabilistic data: Foundations and challenges. In *PODS*, 2007.

[10] O. Dekel, C. Manning, and Y. Singer. Log-linear models for label-ranking. In *NIPS 16*, 2004.

[11] A. Deshpande, C. Guestrin, and S. Madden. Using probabilistic models for data management in acquisitional environments. In *CIDR*, 2005.

[12] Xin Luna Dong, Alon Y. Halevy, and Cong Yu. Data integration with uncertainty. In *VLDB*, 2007.

[13] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation methods for the web. In *WWW*, 2001.

[14] Ronald Fagin, Ravi Kumar, and D. Sivakumar. Comparing top k lists. In *SODA*, 2003.

[15] Jensen Finn and Jensen Frank. Optimal junction trees. In *UAI*, pages 360–366, 1994.

[16] N. Fuhr and T. Rolleke. A probabilistic relational algebra for the integration of information retrieval and database systems. *ACM Trans. on Info. Syst.*, 1997.

[17] Minos Garofalakis and Dan Suciu, editors. *IEEE Data Engineering Bulletin Special Issue on Probabilistic Data Management*. March 2006.

[18] Ronald Graham, Donald Knuth, and Oren Patashnik. *Concrete Mathematics: A Foundation for Computer Science*. Addison-Wesley, 2 edition, 1994.

[19] Todd Green, Grigoris Karvounarakis, and Val Tannen. Provenance semirings. In *PODS*, pages 31–40, 2007.

[20] Todd Green and Val Tannen. Models for incomplete and probabilistic information. In *EDBT*, 2006.

[21] R. Gupta and S. Sarawagi. Creating probabilistic databases from information extraction models. In *VLDB*, 2006.

[22] R. Herbrich, T. Graepel, P. Bollmann-Sdorra, and K. Obermayer. Learning preference relations for information retrieval. In *ICML-98 Workshop: Text Categorization and Machine Learning*, page 8084, 1998.

[23] M. Hua, J. Pei, W. Zhang, and X. Lin. Ranking queries on uncertain data: A probabilistic threshold approach. In *SIGMOD*, 2008.

[24] I. Ilyas, G. Beskales, and M. Soliman. A survey of top-k query processing techniques in relational database systems. *ACM Computing Surveys*, 2008.

[25] T. S. Jayram, R. Krishnamurthy, S. Raghavan, S. Vaithyanathan, and H. Zhu. Avatar information extraction system. *IEEE Data Eng. Bull.*, 29(1), 2006.

[26] J.F.Hauer, C.J. Demeure, and L.L.Scharf. Initial results in prony analysis of power system response signals. *IEEE Transactions on Power Systems*, 5(1):80–89, 1990.

[27] C. Jin, K. Yi, L. Chen, J. Xu Yu, and X. Lin. Sliding-window top-k queries on uncertain streams. In *VLDB*, 2008.

[28] T. Joachims. Optimizing search engines using click-through data. In *Proc. ACM SIGKDD*, pages 133–142, 2002.

[29] L. Lakshmanan, N. Leone, R. Ross, and V. S. Subrahmanian. Probview: a flexible probabilistic database system. *TODS*, 1997.

[30] Jian Li and Amol Deshpande. Consensus answers for queries over probabilistic databases. PODS, 2009.

[31] C. Re, N. Dalvi, and D. Suciu. Efficient top-k query evaluation on probabilistic data. In *ICDE*, 2007.

[32] A. Sarma, O. Benjelloun, A. Halevy, and J. Widom. Working models for uncertain data. In *ICDE*, 2006.

[33] P. Sen and A. Deshpande. Representing and querying correlated tuples in probabilistic databases. In *ICDE*, 2007.

[34] M. Soliman, I. Ilyas, and K. C. Chang. Top-k query processing in uncertain databases. In *ICDE*, 2007.

[35] J. Widom. Trio: A system for integrated management of data, accuracy, and lineage. In *CIDR*, 2005.

[36] K. Yi, F. Li, D. Srivastava, and G. Kollios. Efficient processing of top-k queries in uncertain databases. In *ICDE*, 2008.

[37] Xi Zhang and Jan Chomicki. On the semantics and evaluation of top-k queries in probabilistic databases. In *DBRank*, 2008.

A Multiplication of a Set of Polynomials

Given a set of polynomials $P_i = \sum_{j\geq 0} c_{ij}x^j$ for $1 \leq i \leq k$, we want to compute the multiplication $P = \prod_{i=1}^{k} P_i$ written in the form $P = \sum_{j\geq 0} c_j x^j$. Let $d(P)$ be the degree of the polynomial P.

First we note that the naive method (multiply P_is one by one) gives us a $O(n^2)$ time algorithm by simple counting argument. Let $\bar{P}_i = \prod_{j=1}^{i} P_j$. It is easy to see $d(\bar{P}_i) = \sum_{j=1}^{i} d(P_i)$. So the time to multiply $\bar{P}_i$ and P_{i+1} is $O(d(\bar{P}_i) \cdot d(P_{i+1}))$. Then, we can see the total time complexity is

$$\sum_{i=1}^{k-1} O(d(\bar{P}_i) \cdot d(P_{i+1})) = O(n) \cdot \sum_{i=1}^{k-1} d(P_{i+1}) = O(n^2).$$

Now, we show how to use divide-and-conquer and FFT(Fast Fourier Transformation) to achieve a $O(n \log^2 n)$ time algorithm. It is well known that the multiplication of two polynomial of degree $O(n)$ can be done in $O(n \log n)$ time using FFT. The divide-and-conquer algorithm is as follows: If there exists any P_i such that $d(P_i) \geq \frac{1}{3}d(P)$, we evaluate $\prod_{j:j\neq i} P_i$ recursively and then multiply it with P_i using FFT. If not, we partition all P_is into two sets S_1 and S_2 such that $\frac{1}{3}d(P) \leq d(\prod_{i\in S_i} P_i) \leq \frac{2}{3}d(P)$. Then we evaluate S_1 and S_2 separately and multiply them together using FFT. It is easy to see the time complexity of the algorithm running on input size n satisfies

$$T(n) \leq \max\{T(\frac{2}{3}n) + O(n \log n), T(n_1) + T(n_2) + O(n \log n)\}$$

where $n_1 + n_2 = n$ and $\frac{1}{3}n \leq n_1 \leq n_2 \leq \frac{2}{3}n$. By solving the above recursive formula, we know $T(n) = O(n \log^2 n)$.